Payment Matters? – An Exploratory Study into Pre-Payment Electricity Metering

Philipp-Bastian Brutscher

February 2011

CWPE 1124 & EPRG 1108

EPRG WORKING PAPER

UNIVERSITY OF CAMBRIDGE

Electricity Policy Research Group

Payment Matters? - An Exploratory Study into Pre-Payment Electricity Metering

EPRG Working Paper 1108

Cambridge Working Paper in Economics 1124

Philipp-Bastian Brutscher

Abstract

In this paper we look at the role of pre-payment (in the context of pre-payment metering) for household electricity consumption. Using a matching approach, we find that households paying their electricity up-front tend to consume no less electricity than households paying ex post. This is despite facing a higher tariff and higher transaction costs.In the second part of the paper, we explore to what extent this finding can be linked to an increase in payment flexibility under a pre-payment regime. Using data from the main electricity supplier in Northern Ireland (NIE Energy), we explore how people top-up their pre-payment meters and whether there is a link between people's top-up behaviour and their electricity consumption.

Keywords Pre-payment Metering, Top-Up Behaviour, Demand Management, Pain of Paying.

JEL Classification D03, D04, D12

Contact pbb24@cam.ac.uk
Publication February 2011
Financial Support SuperGen, Flexnet

E·S·R·C ECONOMIC & SOCIAL RESEARCH COUNCIL

www.eprg.group.cam.ac.uk

Payment Matters? - An Exploratory Study into Pre-Payment Electricity Metering

Philipp-Bastian Brutscher*
Electricity Policy Research Group and Faculty of Economics
University of Cambridge

February 10, 2011

Abstract

In this paper we look at the role of pre-payment (in the context of pre-payment metering) for household electricity consumption. Using a matching approach, we find that households paying their electricity upfront tend to consume more electricity than households paying ex post. This is despite receiving information feedback on their electricity use and facing higher transaction costs.

In the second part of the paper, we explore to what extent this finding can be linked to an increase in payment flexibility under a pre-payment regime. Using data from the main electricity supplier in Northern Ireland (NIE Energy), we explore how people top-up their pre-payment meters and whether there is a link between people's top-up behaviour and their electricity consumption.

1 Introduction

With improvements in technology and falling operating costs pre-payment electricity metering is experiencing a revival all across Europe. A particularly impressive example is the case of Northern Ireland: To date more than 240,000

*I would like to thank my supervisor, David Newbery, for reading and discussing numerous drafts of this paper. I would also like to thank Michelle Baddeley, Tom Crossley, Jonathan Grant, Hamish Low, Bhanu Patruni, Michael Pollitt and an anonymous reviewer for commenting on and helping to improve the paper at various stages of its development. Finally, I would like to thank Sarah Beldon, Jenny Boyd and Graeme Hunter from NIE Energy for their interest in the project and their great support. Financial support from the EPSRC (Flexnet) is gratefully acknowledged. The usual disclaimer applies.

households in Northern Ireland (ca 30%) use pre-payment metering to pay for their electricity.

Despite the wide-spread uptake, little is known about the role of pre-payment for household energy consumption. This lack of research is surprising: Clearly, household energy consumption is high up on the policy agenda as reflected in a series of high-level policy reports – including the NAO's *Programmes to Reduce Household Energy Consumption* and the Carbon Trust's report *Climate Change: a Business Revolution*.

In addition, a large body of literature suggests that 'payment matters' for consumer behaviour: People's consumption behaviour has been shown to depend on the payment method (Hirschman, 1979; Prelec and Simester, 1998); the time between payments (Gourville and Soman, 1998); the way payments are 'framed' (Gourville, 1998); and the extent to which payments are "bundled" (Morwitz et al, 1998; Chetty et al, 2010).

In this paper we make a first step towards better understanding the role of payment in the context of pre-payment electricity metering. Using data from the Northern Ireland Continuous Household Survey; from the main electricity provider in Northern Ireland (NIE Energy); and the Northern Ireland Neighbourhood Information Service, we focus on three questions:

1. What is the effect of pre-payment on household energy consumption compared to post consumption payment?

2. How do consumers use their pre-payment meters – e.g. what payment schedules do they choose? and

3. What is the relationship between how consumers use their meters and their energy consumption – e.g. does purchasing smaller top-ups more often make people consume more energy?

The paper is organised as follows: In the first part, we briefly describe the pre-payment situation in Northern Ireland. In the second part, we evaluate the effect of pre-payment on electricity consumption (relative to post-consumption payment). The third part looks at people's payment schedule under a pre-payment system – highlighting two behavioural anomalies. In the fourth part, we explore the link between households' electricity use and their payment schedules. We conclude the paper with a discussion of some preliminary policy implications arising from our research.

2 Pre-payment metering in Northern Ireland

We start our discussion with some background information on pre-payment metering in Northern Ireland.

2.1 Background

Today customers can devise to pay their electricity bills in a number of ways. They can pay cash or by cheque, use direct debit, or paperless online billing (where customers take their own meter readings and enter them online). Another way of paying for electricity, which is gaining increasing popularity, is by means of pre-payment metering.

The general idea of pre-payment metering is that electricity can only be consumed if one's meter is in credit. When credit runs out, electricity supply is stopped. While traditionally pre-payment meters were used primarily in tenements buildings and individually rented rooms, today they are used across all socio-economic groups: To date, about 30% of all electricity customers in Northern Ireland use pre-payment metering – with new connections continuing at a rate of 2,000 per month.

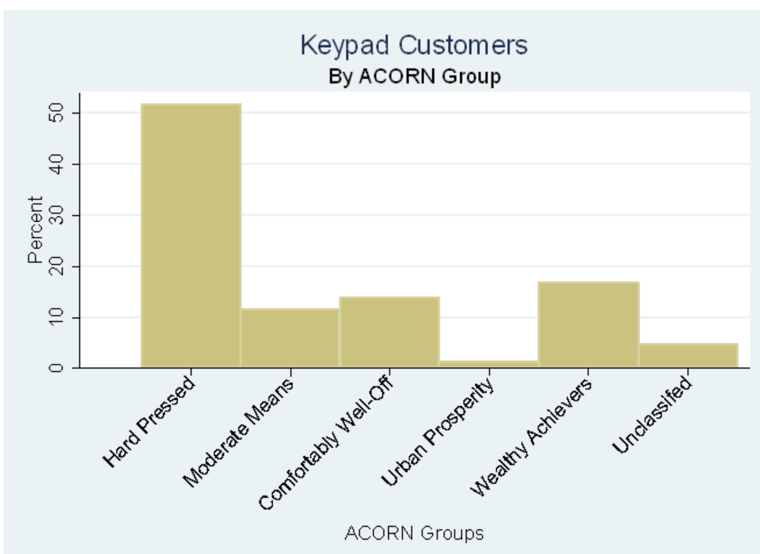

Figure 1: Keypad Customers by ACORN group

Figure 1 shows the distribution of pre-payment customers in Northern Ireland by ACORN group. ACORN groups are defined on the basis of several socio-economic variables.[1] For simplicity, we focus here on income. The figure shows that while half of all pre-payment customers fall into the ACORN group 'hard pressed' – with an average family income of 60% of the UK mean – 35% fall into the highest ACORN groups 'comfortably well-off' (101% of mean income), 'urban prosperity' (129% of mean income), and 'wealthy achievers' (137% of mean income).

2.2 Main Features of the Keypad meters

The widespread uptake of pre-payment metering in Northern Ireland is closely linked to a change in technology: Being dissatisfied with the old system, in 2002, NIE Energy switched from a smart card system to a keypad metering system. Only after this change in technology, pre-payment metering became broadly used.[2]

The keypad system works in a similar way to a mobile phone top-up system. At the vending outlet customers purchase a 16-digit code which they enter into the keypad of their meter to receive credit. Other codes can be issued/used to recalibrate the meter or to change the settings of the meter.[3] From a customer's perspective, the main attraction of the new keypad system is that keypad meters come with a discount compared to standard credit of 2.5% - which compares to a discount of 4% for direct debit customers.

In addition, the keypad meter has a conveniently placed display which enables customers to monitor consumption, credit available etc. It is also felt that the range of credit top-up facilities[4] – customers can purchase their top-up in shops (Payzone; Paypoint); post offices; via the phone; or using the internet – have attracted a broader range of users and helped remove the stigma of pre-payment.[5]

---

[1] ACORN is a geo-demographic information system categorising all United Kingdom postcodes into various types based upon census data and other information such as lifestyle surveys.

[2] See Zhang (2010) for a detailed discussion of the diffusion process.

[3] There is no two-way communication – which is why the meter is typically referred to as 'semi-smart'.

[4] The majority of top-ups are (still) purchased at Paypoints and Payzone outlets. However, this picture varies for different top-up amounts – with relatively more customers using phone or internet top-ups at higher top-up amounts. In addition, internet and phone top-ups increased by 37% in 2009 compeared to 2008.

[5] A further attraction of the new pre-payment system is that it comes with 'friendly credit'. This means: users cannot self-disconnect at weekends or between 4pm and 8am (which can be extended to 11am on request). This safeguard was requested by Ofreg, the Northern Ireland

The main advantage of the new system for NIE Energy is that keypad meters come with reduced management costs: they do not require manned meter readings, disconnections, or re-connections. Call outs occur only during working hours and the danger of inaccurate meter readings is essentially eliminated. In addition, there is no need to issue bills and the handling of debts becomes much easier: Every time a customer in debt buys credit, a fixed fraction of his/her top-up can be marked towards redemption of old debts.

2.3 Pre-payment customers in Northern Ireland

In the last section, we noted the uptake of keypad meters across all socio-economic groups. For what follows, it is important to bear in mind, however, that despite the wide-spread acceptance of the new technology, pre-payment customers are not a representative sample of the population.

The 2008/2009 Northern Ireland Continuous Household Survey asked 2,632 households about their electricity consumption and how they pay for it. This is in addition to a large number of questions about their background characteristics and other consumption behaviour. This information allows us to draw out some of the differences between keypad customers and customers using other forms of payment for their electricity.

Table 1 and 2 provide summary statistics of the data. Table 1 shows the distribution of electricity customers in Northern Ireland by payment type. Keypad customers fall into the group 'Slot-meter, Power Card & Pay-As-You-Go'. Because NIE Energy was the only electricity supplier in the residential sector at the time, and because it provides only one type of pre-payment meter (its keypad meter), this group comprises keypad customers only.

regulator, due to concerns from consumer groups and others about self-disconnection. Clearly, electricity used during periods of friendly credit has to be repaid at the next top-up (Owen and Ward, 2010).

	Number of Observations
Account	483
Monthly Instalments	594
Budget Account	62
Slot, Power Card and Pay-As-You-Go	746
DHSS Direct Payment	33
Cash/Check with bill	691
Total	2,609

Table 1: Distribution of electricity customers in NI

The table shows: customers paying their electricity bill in cash or by check and cutomers paying by means of pre-payment make up the largest portion of customers – which is consistent with the situation for the whole of Northern Ireland. Table 2 below shows the mean values of a series of background characteristics for the entire sample.

	All households
Age of HH head	42.08
	(16.35)
Female HH head	0.47
	(0.50)
HH head eco. inactive	0.40
	(0.49)
HH income (£1,000)	14,969
	(15,866)
Number of adults	1.90
	(1.00)
Number of children	0.67
	(1.01)
Detached	0.12
	(0.33)
Number of rooms	4.97
	(1.16)
Old building	0.04
	(0.20)
Renting	0.61
	(0.49)
Electric Heating	0.26
	(0.44)

Table 2: Background variables Keypad Cusomters and Others

In Table 3, we report the results from a logistic analysis on the association between household background characteristics and the use of a keypad meter. Economic theory provides little guidance on the specification of the model, so variables were chosen on an ad hoc basis. The left hand column of Table 3 provides the estimated coefficients; the right hand column shows the marginal effect of each variable on the probability of having a keypad meter calculated at the mean of the explanatory variable.

	Coefficient	Marginal Effect
Age of HH head	0.05***	0.008***
	(0.02)	0.003
Age of HH head squared	-0.0009***	-0.0001***
	(0.0002)	0.00003
Female HH head	0.08	0.01
	(0.11)	0.02
HH head eco. inactive	0.60***	0.10***
	(0.14)	0.02
HH income (£1,000)	-0.02***	-0.003***
	(0.003)	0.001
Number of adults	0.086	0.01
	(0.058)	0.01
Number of children	0.18***	0.03***
	(0.06)	0.01
Detached	-1.22***	0.18***
	(0.14)	0.02
Number of rooms	-0.16***	-0.03***
	(0.04)	0.01
Old building	-0.35	-0.05*
	(0.23)	0.03
Electric Heating	-0.04	-0.006
	(0.12)	(0.02)
Renting	0.84***	0.15***
	(0.12)	0.02
Constant	-0.23	
	(0.48)	
Log likelihood	-1181.95	P(1-P)=0.17

Table 3: Results from Logistic Analysis.
Statistically Significant at 1% ***; 5% **; 10% *.

The results from our estimation show that, all else equal, keypad customers tend to be older, have lower average incomes, and are more likely to be economically inactive. In addition, we find that keypad customers tend to live in smaller houses – sharing them with more people. We also find that households renting are more likely to have a keypad meter than households owning a property/flat.

In terms of absolute size of the effects, we find that – at the mean values of the variables – an additional life year of the household head is associated with a 0.8 percentage point increase in the probability of uptake of a keypad meter. An increase in property size by one room, on the other hand, is associated with a decrease in the probability of having a keypad meter of 3 percentage points.

These findings illustrate that despite the widespread uptake of keypad meters, the differences between keypad and non-keypad customers are (still) significant.

3 The Effect of Pre-payment

In this section, we evaluate the effect of using a keypad meter on household electricity consumption.

3.1 Related Research

There is a large body of evidence showing that information feedback (on electricity consumption) typically leads to a decrease in household electricity use. The research is summarised in Table 4 below.[6] At the same time, little is known about the effect of pre-payment on household electricity use.

[6] A recent qualitative study comes from Hargreaves, Nye and Burgess (2010).

Study	Type of Information	Results	Comments
Darby, 2006	Direct Feedback: - Self-meter reading - Direct displays - Interactive feedback	5 to 15% savings	Range of international studies with different types of direct feedback
Darby, 2006	Indirect Feedback: - Frequent bills - Frequent bills based on readings plus other historical/comparative/detailed information	0 to 10% savings	Range of international studies with different types of indirect feedback
Wood and Newborough, 2003	Electronic feedback via consumption indicator attached to electric cooker.	15% reduction	44 UK households; focuson electricity for cooking
Wood and Newborough, 2003	Paper-based information pack on electricity consumption of cooking appliances and electricity savings tips	3% reduction	
Dulleck and Kaumann, 2004	Information leaflets on energy efficiency; introd. of energy efficiency appliance certi.	7% reduction	Impact on long-run rather than short-run demand

Table 4: Literature on Information Feedback. Table adopted and adapted from Brophy Haney et al (2009)

An initial trial of 200 households in Northern Ireland found an average of 10% electricity saving with keypad meters (NIE Energy, 2000). However, the households in the trial which received a new pre-payment meter were hand-held during the process.[7] This is likely to have affected their consumption behaviour - which makes it hard to judge the impact estimate in terms of its external validity.

[7] The group of households was a convenience sample. In addition, they were contacted regularly to provide feedback on their experience.

Follow up research on a broader sample found an average of 3% saving (NIE Energy, 2003). Yet, the evaluation involved a time of day tariff. This gave customers a strong incentive to save at peak times, which, again makes it hard to draw firm conclusions about the effect of pre-payment/the keypad on household electricity use.[8]

The lack of research on the effect of pre-payment metering is surprising: understanding whether and to what extent pre-payment (too) affects household energy consumption is important for our understanding of household consumption behaviour. In addition, as pointed out by Fischer (2008), only if we understand the mechanisms that influence household behaviour will we be able to manage energy preservation effectively.

3.2 A Matching Approach

In the following, we estimate the effect of having a keypad meter on electricity consumption.

A naive way of assessing the effect of the keypad is by comparing the electricity use of households with a keypad meter and households without it. It becomes clear very quickly, however, that this is uninformative: since the two groups of households are very different from each other, any difference in electricity use is likely to reflect not only the effect of having/not having a keypad but also differences in income, housing, living arrangements etc.

What we need to know to evaluate the effect of the keypad is what electricity consumption of households with a keypad meter would have been, had they not had a keypad meter. That is, what we need to know is the counter-factual. The evaluation problem arises, because we do not observe this counterfactual. All we observe is the electricity consumption of households with and without a keypad meter.

The recent evaluation literature has focused on matching estimators to overcome this problem. (See Dehejia and Wahba, 1999 and Heckman et al, 1998). The basic idea of matching (applied to our context) is that the bias in evaluating the effect of the keypad meter on electricity consumption is reduced when the comparison of consumption is performed using households which are as similar as possible.

[8] The literature sometimes refers to the 'M Power Conservation Effect Study' (Pruitt, B., 2004) and the Woddstock Hydro evaluation (reported in Quesnelle, K., 2004). However, we could not find much information on these projects - in particular little is known about the method(s) which were used in these studies.

3.3 Matching Formally

Suppose we have data on T keypad customers and C non-keypad customers. In addition, suppose we have a vector X of variables which help predict whether or not a household has a keypad meter. Given this data, we can match keypad customers with a comparison group of (similar) non-keypad customers.

Ideally, we would match each customer using a keypad meter with a single non-keypad customer that has an identical value of X. This is unpractical, however, because the dimension of X might be high: as the number of characteristics used in matching increases, the chance of finding an exact match becomes smaller and smaller.

Rosenbaum and Rubin (1983) show that matching can also be performed using P(X) rather than X – where P(X) is the probability of having a keypad meter conditional on X, i.e. the 'propensity score'.[9] The propensity score can be calculated for each household using standard discrete choice parametric or semi-parametric models. We use standard parametric likelihood methods to compute the propensity score.[10]

We then use the odds ratio $p_i = P_i/(1-P_i)$ where P_i is the estimated probability for a household to have a keypad meter, to construct matched pairs.[11] Because it is unlikely that two individuals have the exact same score, several matching algorithms have been suggested in the literature (Becker and Ichino, 2002). We focus on the two most popular ones: nearest neighbour matching and kernel matching.

The nearest neighbour to the i-th household is defined as the non-keypad household which minimises $[p(X_i)-p(X_j)]$^2 over all j households in the set of non-keypad households – where $p(X_n)$ is the predicted odds ratio for observation n. Sometimes, nearest neighbours may (still) be far apart in terms of the distance metric between the propensity scores of households with a keypad meter and those without.

This is why we also use a kernel estimator: it takes into account all available information and puts a greater weight on good matches than on bad ones. That is, it matches electricity consumption of households with a keypad meter with that of all non-keypad households, with weights that are inversely proportional

[9] More specifically, they show that if potential outcomes are independent of treatment conditional on the available characteristics, they are also independent of treatment conditional on the propensity score.

[10] Several studies show that the impact estimator is robust to the choice of the discrete choice model (see e.g. Heckman et al, 1998).

[11] We allow for replacement.

to the distance between the propensity scores of the keypad and non-keypad households.

The mean impact estimator (τ) becomes:

$$\tau^k = \frac{1}{T}\sum_{j\in C(i)}[Y_i^T - \frac{\sum_{j\in C} Y_j^C G(\frac{p_j - p_i}{h_n})}{\sum_{k\in C} G(\frac{p_k - p_i}{h_n})}] \quad (1)$$

where the last term can be interpreted as an estimator of the counterfactual. Y^T and Y^C are the observed outcomes in terms of electricity consumption of households with a keypad meter, households without it, respectively. G(.) is a kernel function – in our case the Epanechnikov kernel – and h_n a bandwidth parameter.[12]

3.4 A Note on the data

For our estimation, we use data from the 2008/2009 Northern Ireland Continuous Household Survey. In order to get a more meaningful comparison group, we exclude all but account paying households and households paying their electricity bill by cash or check from the group of non-keypad households.[13] In addition, we drop households with electriciy consumption less than 1 KWh/day.[14]

In Table 5, we present selected descriptive statistics for keypad and non-keypad households.[15]

[12] In order to increase the precision of our estimates we allow for replacement - i.e. the use of the same comparison household for several treatment households.

[13] We were worried that expenditure information on customers paying by monthly instalments might not reflect their actual consumption. Similarly, we were concerned with the marginal costs customers face for whom electricity is paid directly by DHSS. See Borenstein (2009)

[14] A refrigerator typically uses 1-2 KWh/day, so it is implausible that an occupied primary residence woul would fall below 1 KWh/day.

[15] Please note: electricity use is calculated assuming a 2.5% discount for keypad customers and a 4% discount for account paying customers. Please note also: the CHS data explicitly excludes arrears (or rental charges).

	PPM	Account/Cash&Check
Number of Observations	746	1,174
Electricity Use (KWh/Qrtly)	753.3	786.8
	(635.6)	(533.8)
Age of HH head	42.08	55.36
	(16.35)	(18.52)
Female HH head	0.47	0.35
	(0.49)	(0.47)
HH head eco. inactive	0.40	0.40
	(0.49)	(0.49)
HH income (£1,000)	14,969	21,391
	(15,866)	(23,703)
Number of adults	1.90	2.01
	(1.00)	(0.96)
Number of children	.67	0.46
	(1.00)	(0.93)
Detached	0.12	0.47
	(0.33)	(0.50)
Number of rooms	4.97	5.93
	(1.16)	(1.96)
Old building	0.04	0.11
	(0.20)	(0.31)
Electric Heating	0.26	0.23
	(0.43)	(0.42)
Renting	0.61	0.23
	(0.49)	(0.43)

Table 5: Descriptive Statistics of Keypad and Non-keypad customers in NI

3.5 Estimation

Estimating the propensity score is a crucial step in using matching as an evaluation strategy. Different strategies have been adopted to choose a suitable specification of the treatment equation (see e.g. Dehejia and Wahba, 1998; Smith and Todd et al, 1998). The underlying principle is that 'pre-intervention variables' should be included in the regression which are not influenced by whether or not a household has a keypad meter.

The Continuous Household Survey provides a rich set of information on household and housing characteristics. We estimate alternative logistic models to predict whether a household has a keypad meter and select a final model on the basis of the likelihood function. Table 6 presents the logit regression used to estimate the propensity score on the basis of which the matching is subsequently done.

Variable	Coefficient
Age of HH head	0.05**
	0.02
Age of HH head squared	-0.009***
	0.0001
HH head eco. inactive	0.65***
	(0.15)
HH income (£1,000)	-0.01***
	0.003
Number of adults	-0.04
	0.06
Number of children	0.16***
	0.06
Number of rooms	-0.29***
	0.05
Old building	
Electric Heating	-0.05
	0.13
Renting	0.83***
	0.13
Constant	0.95***
	0.51

Table 6: Propensity Score Estimation.

Statistically Significant at 1% ***; 5% **; 10% *.

After estimating the propensity score for households with and without a keypad meter, we plotted them to check the common support condition (see Lechner, 2000).[16] The plots are shown in Figures 2 below. They show a good overlap in propensity scores between households with and without a keypad meter. We exclude households for the small area for which there is no overlap:

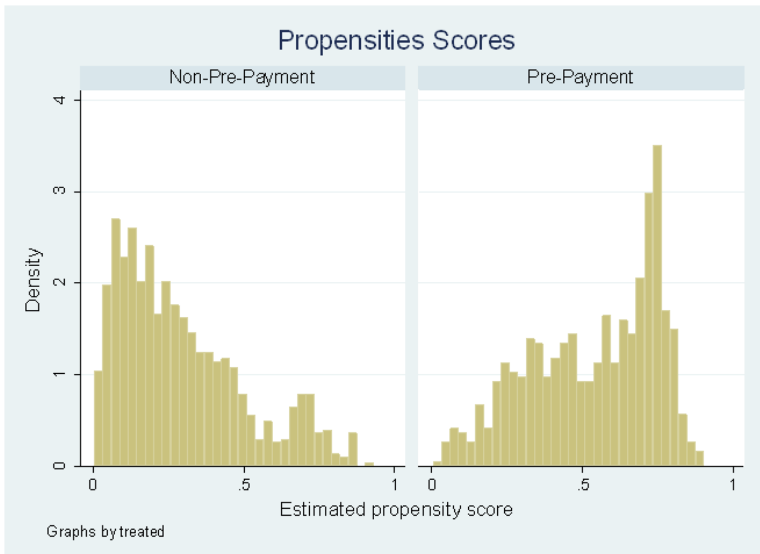

Figure 2: Propensity Scores

Table 7 below gives our estimate of the difference in consumption between keypad customers and account/cash/check paying customers. Our estimate suggests a difference of between 30 and 75 KWh per quarter.

	Nearest Neighbour	Kernerl Matching
Average Treatment Effect	31.70	73.76
on the Treated	(40.68)	(35.84)

Table 7: Average Treatment Effect from Matching Estimator

That is, our findings suggest that households with a keypad meter tend to consume between 30 and 75 KWh more electricity than comparable account/cash/check paying households.

[16] The balancing hypothesis is not satisfied for all variables in the highest block (block 9). To check the robustness of our results, we re-estimate the effect of having a keypad meter using only observations in the region of 'thick support'.

This finding is robust to different specifications of the propensity score and different matching algorithms. In addition, we re-estimated the effect of having a keypad meter using only households in the region of 'thick support' (with a propensity score between 0.20 and 0.80). What we find is that the thick support estimates are very similar to our baseline results.[17]

Finding a higher electricity use for keypad customers is suprising: On the one hand, keypad meters come with a 2.5% discount compared to standard credit. On the other hand, this compares to a 4% discount for direct debit customers.

In addition, even if we abstract from direct debit - given the relatively low price elasiticy of electricity[18] - it is very unlikely that the lower tariff for keypads can explain the entire difference (of between X and Y percent) in electricity use between keypad and non-keypad customers.

What is more: keypad meters come with information feedback on electricity use. In addition, the meters are associated with higher transaction costs: every time a customer needs additional credit, he/she has to travel to an outlet.[19] Both of these aspects suggest that electricity consumption should be lower (rathert than higher) for keypad customers.

3.6 Unobserved Characteristics

One possible explanation for our finding is that one of our key identification assumption does not hold. The assumption is that conditional on all observed household characteristics, whether a household uses a keypad meter or not is independent of its electricity use without a keypad.

One reason why our identification assumption may not hold is that there is some unobservable component which affects both – whether a household uses a keypad meter and how much electricity it consumes. This could be for example the (unobserved) ability to manage one's electricity use.

If households which are worse in managing their electricity use are both – more likely to use a keypad meter and more likely to have a higher electricity consumption in the absence of a keypad meter, then our matching estimator may find no negative (or even a positive) effect of having a keypad meter on electricity use even if it does have a negative effect.

[17] Specifically, what we find is an ATT of 64.1 using the nearest neighbour algorithm and an ATT of 70.0 using our kernel algorithm.

[18] Price elasticity for electricity is typically estimated significantly below 1. See e.g. Lijesen (2007); Halvorsen and Larsen (2001)

[19] The extent of these transaction costs will become clear(er) in the next section.

The problem is: by definition, we cannot test the effect of an unobserved component on our estimation result. That is, we cannot test whether our estimation results are biased. What we can do, however, is test how much an unobserved component (like the ability to manage one's electricity use) would have to influence a household's decision to use a keypad meter and electricity consumption to significantly alter our estimation results.

The basic idea is that, if we can show that all configurations of the unbserved component which lead to a reversal of our findings can be considered unlikely, we can be reasonably sure that our estimates reflect the true effect of using a keypad meter on electricity use (rather than the combined effect of using a keypad meter and, say, being bad in managing one's electricity use).

3.7 Robustness Test

The main steps of our test can be summarised as follows: We first make different assumptions about the distribution of the unobserved component (U) in our sample. We then test under which assumptions of U, the estimated effect of using a keypad meter on electricity use becomes negative. Finally, we discuss how plausible these assumptions are. The approach was first suggested by Ichino et al (2008).

To implement the test, we assume that U is binary and iid distributed. This allows us – without loss of generality[20] – to characterize the distribution of U by means of four parameters:

- P_{11}=Pr(U=1|T=1, Y=1, X)
- P_{10}=Pr(U=1|T=1, Y=0, X)
- P_{01}=Pr(U=1|T=0, Y=1, X)
- P_{00}=Pr(U=1|T=0, Y=0, X)

which are the probabilities that U=1 in each of the 4 groups defined by whether households have a keypad meter or not and whether households have above mean electricity use or not.[21]

To test the effect of a given distribution of U, we fix p_{11}-p_{00} at a given value at a time. We then use these values to attribute a value of U to each individual

[20] Ichino et al (2007) provide Monte Carlo simulations which show that the assumption does not critically affect the results of the robustness test.

[21] We also used a binary transformation using the median. Using the median does not affect the result of the sensitivity analysis).

in our sample. To give an example: If we fix p_{11} at 0.6, we attribute a value of U=1 with probability 0.6 to each household with a keypad meter (T=1) and which has above electricity use (Y=1).

We then estimate the effect of having a keypad meter on electricity use – including U as a further observed covariate. We repeat this process a large number of times for our given set of values of p_{11}-p_{00}. We obtain an estimate of the effect of having a keypad meter, as the average of the estimated effects over the distribution of the simulated Us.

In Table 8 below, we show the findings from this exercise. To reduce the dimensionality problem of the characterization of U, we fixed the probabilities Pr(U=1) and the difference d'=p_{11}-p_{10} at some pre-determined values. This allows us – again without loss of generality[22] – to fully describe the simulated confounder by the difference d=p_{01}-p_{00} and s=$p_{1.}$-$p_{0.}$.[23]

	s=0.1 $\Lambda \in$(1.53; 1.69)	s=0.2 $\Lambda \in$(2.38; 2.64)	s=0.3 $\Lambda \in$(3.68; 4.20)	s=0.4 $\Lambda \in$(5.80; 6.58)	s=0.5 $\Lambda \in$(9.68; 11.68)
d=0.1	71.1	66.5	59.1	46.2	32.4
$\Gamma \in$(1.54;1.84)	(4.0)	(7.1)	(13.8)	(17.6)	(20.9)
d=0.2	66.3	55.8	39.3	17.5	-7.9
$\Gamma \in$(2.38;3.33)	(4.6)	(7.2)	(11.7)	(19.6)	(23.8)
d=0.3	61.3	41.6	-16.6	-17.2	-60.8
$\Gamma \in$(3.90;6.54)	(4.8)	(9.6)	(13.1)	(21.4)	(23.6)
d=0.4	54.3	31.2	-5.4	-47.8	-113.6
$\Gamma \in$(5.98;15.07)	(5.4)	(11.0)	(16.7)	(20.4)	(26.7)
d=0.5	48.0	15.4	-32.2	-92.6	-188.9
$\Gamma \in$(10.30;85.57)	(7.9)	(12.8)	(15.8)	(21.8)	(25.7)

Table 8: Average Treatment Effect from Matching Estimator with an Unobserved Component

The table also shows the estimated values of Λ and Γ – where Λ represents a measure of the effect of U on the probability that a household uses a keypad meter (selection effect) and Γ a measure of the effect of U on electricity use (outcome effect).

[22] Since these quantities are not expected to represent a real threat for the baseline estimate, they can be held fixed and the simulated confounder U can be fully described by the difference d= p_{01}-p_{00} and s=$p_{1.}$-$p_{0.}$.

[23] We only use look at d>0 and s>0. By assuming p_{01}>p_{00} one can simulate a confounding factor that has a positive effect on the untreated outcome y0. Similarly, by setting $p_{1.}$>$p_{0.}$ One can simulate a confounding factor that has a positive effect on treatment assignment.

The key finding from the table is that even if we allow U to be distributed in a way that it has a large effect on the probability that a household uses a keypad meter (Λ=7.6) and a large effect on electricity use (Γ=3.8), the estimate of the corresponding effect of using a keypad meter on electricity use remains positive.[24]

To reverse our (baseline) estimate to give a negative effect of about 70 KWh, U needs to have a very (and implausibly) large effect on the probability that a household uses a keypad meter and/or its electricity use. More specifically, U needs to increase the relative probability of using a keypad meter (Λ) by a factor greater than 10.3 (6.6) and the relative probability of having above average electricity consumption (Γ) by a factor greater than 6.5.(23.2).

For comparison, if we model U to resemble our 'electric heating' dummy (a dummy indicating whether a household has at least one child), we get a selection of Λ=1.23 (2.09) and an outcome effect of Γ=1.07 (3.44). This suggests that it is unlikely that our finding that using a keypad meter tends to increase electricity consumption is driven by some unobserved component.

4 Exploring people's top-up behaviour

In the last section, we found that having a keypad meter tends to increase electricity consumption. In this section, we describe how people use their keypad meters: we look at how often they top-up their meters and how much money they typically put on their meter.

What we are interested is to find a cue/anomaly in people's behaviour which can help us explain the positive effect of having a keypad meter on household electricity consumption. We test for such a link more rigorously in the final section of the paper.

For our analysis we use data from NIE Energy on 10,124 randomly chosen households.[25] This corresponds to roughly 2.4% of all keypad customers in Northern Ireland. Our dataset tracks households over 18 months (between 1.6.2008 and 30.11.2009). To have a clear panel data-set, customers who have moved house within this period were not included in the sample.[26]

[24] Please note: the matching uses our kernel algorithm. Standard errors are calculated using the between-imputation variance.

[25] The data was given to us on a confidential basis.

[26] Our data set includes information on when households purchase top-up; what amount they purchase; what channel they use (i.e. online; phone; Call Centre; IVR_Online; Paypoint; Payzone; or Post office); and which tariff they are on. In addition, we can link each household

4.1 A simple model of people's top-up behaviour

From a theoretical perspective, an intuitive way to think about people's top-up beavhiour is in terms of an application of the Baumol Tobin model (Baumol, 1952; Tobin, 1956; Romer, 1986). It suggests that people trade-off the costs and benefits of putting money on their meter. The main benefit is convenience: people put money on their meter to avoid having to purchase top-up every time they need electricity. The cost of this convenience is forgone interest.

To see how people trade-off these benefits and costs, consider a person who spends £Y on electricity over the course of a year. For simplicity, assume that the price level is constant, so real spending is constant over the year. Suppose the individual makes N top-ups and spends £Y/N on each top-up. Suppose that our individual spends the credit she purchases gradually over the next 1/Nth of the year.

In addition, suppose that the cost of topping-up one's meter is some fixed amount F. We can view F as representing the value of time spent to purchase top up (which means to travel to a store; call a helpline; or go online to purchase a vend code before entering it into one's meter at home). Finally, let i denote the interest rate; we can interpret i as a measure of the opportunity cost of holding money on one's meter (since money on a meter does not bear interest).

Now we can analyse the optimal choice of N, which determines the optimal amount of credit to be purchased. For any N, the average amount of money held is Y/2N, so the forgone interest is (Y/2N)i. Because F is the cost per top-up, the cost of top-ups over the course of a year is FN. This means, the total cost our individual bears is:

$$(Y/2N)i + FN \quad (2)$$

which is the sum of the forgone interest and the cost of purchasing top-up.

Simple minimisation gives the optimal value of N and average top-up amount: the optimal value of N is the square root of Yi/2F; the corresponding average top-up amount is the square root of YF/2i. This means that an individual purchases more credit at a time if the fixed cost of purchasing credit, F, are higher; if her expenditure Y is higher; or if the interest rate i is lower.

to a post-code. This allows us to match our data with the Northern Ireland Neighbourhood Information Service (NINIS) database – a locational data set that contains information on a large number of socio-economic variables.

4.2 Emprical puzzle

Two testable implications follow from our model: It suggests that an average customer

- will purchase credit worth £230 about 2.3 times a year and
- will respond to a change in expenditure by changing both: the number of top-ups and average top-up amount.

Take the average annual electricity bill in Northern Ireland – which is £522. Suppose that it takes our individual on average about 10 minutes to top-up her keypad meter;[27] that she has a value of time equal to her average hourly salary of £12;[28] and earns 4 percent annual interest on balances held at her bank.

Plugging this information into our model, we find that, on average, our customer should purchase credit worth £230 about 2.3 times a year and hold an average of £115 worth of credit on her meter.

This contrasts starkly with the behaviour we observe: people top up their meters ca 45 times a year with, on average, £13 at a time. Figure 3 plots the average amounts of credit people purchase every time they top-up their meter for the time period 1.6.2008-30.11.2009 (which includes 2 increases in minimum top-up).

[27] More than 90% of customers top-up their meters in stores. So while it is possible to spend less time on average top-up one's meter using online top-ups, 10 minutes seem a fair estimate for the average time spend to purchase credit.

[28] Source: Northern Ireland Statistics

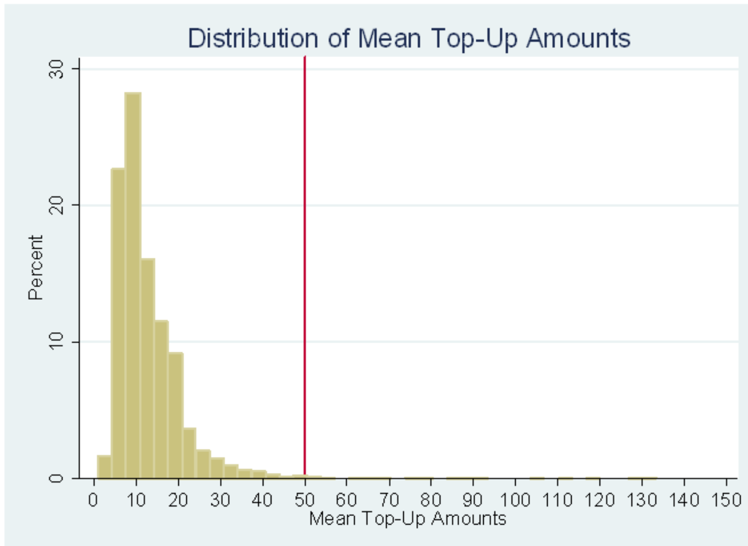

Figure 3: Distribution of Mean Top-up Amounts

The figure shows that the vast majority of customers (96%) top up their meters on average by less than £30 at a time. Only 5 in 1000 purchase credit in excess of £50.

The second testable implication of our model is that a change in expenditure Y should lead to an equal change in the number of top-ups and average top-up amount. (This follows from the formula of the optimal value of N and average top-up amount). Again, the prediction is at odds with the data. Table 9 below shows the change in the number of top-ups and average top-up amount from:

- June to July 2008
- September to October 2008;
- December 2008 to January 2009.
- September 2009 to October 2009

It shows that people respond to increases in tariff mainly through increases in the number of top-up trips and much less through increases in the average top-up amount. The table also shows that this cannot be explained by slow adjustment processes whereby people learn about the increase in tariff over time. We find little adjustment in the two months following the increase in tariff.

	Change: tariff	Change: # of top-ups	Change: average top-up amount
06/08-07/08	0.01588	4,045	£0.01
	(14.0%)	(10.5%)	(0.1%)
07/08-08/08	0	1,489	£0.19
	(0%)	(3.5%)	(2%)
08/08-09/08	0	37	£0.19
	(0%)	(0%)	(1.9%)
09/08-10/08	0.04319	12,428	£0.58
	(33.3%)	(28%)	(5.8%)
10/08-11/08	0	480	£0.44
	(0%)	(0.8%)	(4%)
11/08-12/08	0	-1,740	£0.85
	(0%)	(-3%)	(7.7%)
12/08-01/09	-0.01863	-2,958	-£0.74
	(-10.8%)	(-5.2%)	(-6.2%)
01/09-02/09	0	-1,543	-£0.12
	(0%)	(-2.9%)	(-1%)
02/09-03/09	0	-1,973	-£0.23
	(0%)	(-3.8%)	(-2.1%)
09/09-10/09	-0.768	-1,287	-£0.17
	(-5.0%)	(-2.2%)	(-1.7%)
10/09-11/09	0	784	£0.39
	(0%)	(1.7%)	(3.6%)

Table 9: Overall response to changes in Tariffs

Interestingly, when it comes to decreases in tariffs, we find that, in line with the predictions of our model, people respond to the change in tariff through decreases in average top-up amount as much as through decreases in the number of top-ups.

4.3 Changing the parameters of the model

One possible explanation for the discrepancy between the predictions of our model and the top-up behaviour we observe is the choice of parameters. For example, it may be that average wages measure the opportunity cost of time correctly only if a customer actually works: If he/she is not active in the labour market, as is the case for a significant number of keypad customers, wages may be a rather crude proxy for their valuation of time.[29]

It is also possible that people tend to purchase their top-up when they do their grocery shopping. In this case, arguably, the additional time spent on purchasing top-up is likely to be small and, hence, the marginal (opportunity) cost of doing so low. There are two problems with these arguments: The first one is that while changing the parameters of our model can reduce the difference between the model's predictions and actual top-up behaviour, it cannot explain the full difference.

Figure 4 shows the optimal top-up amount for low and high wage customers (£6, £24 per hour) and low and high interest rates (2% and 6%).

[29] In addition, wages may not capture the true valuation of time even for those who are active in the labour market: To the extent that direct-income compensation for additional time expenditures are generally available only for the self-employed and other autonomous individuals, wages may not reflect most people's valuation of time outside working hours (which is presumably when they purchase top-up for their meters).

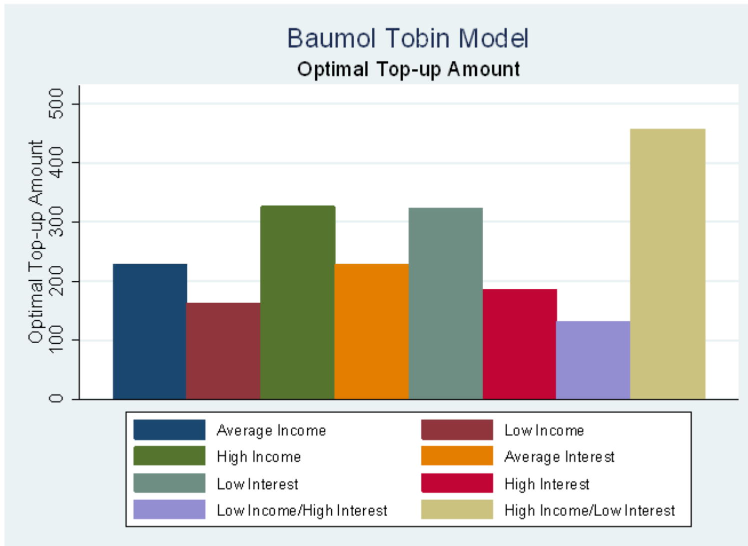

Figure 4: Changing the Parameters of the Model

The figure shows that in no scenario of wages and/or interest rate the optimal top-up amount drops below £130. In fact, even if we assume the most extreme case (not shown) with low income, high interest rates and 1 minute per top-up, we arrive at an optimal top-up amount of £50, which exceeds the average top-up amount of 99% of customers in Northern Ireland.

The second problem with the argument that our anomalies are simply a matter of choosing the right parameters is that customers tend to adjust to increases in tariffs almost exclusively by increasing the number of top-up trips. As outlined above, if our model were an accurate description of how people choose to top-up their meters and we simply had to find the right parameters to fit it, then increases in tariff should result in an equal change in the number of top-up trips and the average top-up amount purchased.

4.4 Liquidity Constraints

An alternative explanation for our two anomalies is liquidity constraints. The idea is that people may be very bad in saving and so make top-ups whenever they have money (rather than after saving for one big top-up). In the light of our earlier discussion and the finding that keypad customers tend to be relatively poorer, liquidity constraints seem to be a realistic explanation for the low top-ups of some customers and the 'inability' to adjust to increases in tariff by increasing (also) the average top-up amount.

However, liquidity constraints cannot explain why also customers in the highest ACORN groups ('comfortably well off', 'urban prosperity', 'wealthy achievers') come nowhere close to the optimal top-up amount suggested by our framework; nor why customers in these groups show the same (unexpected) adjustment process to changes in tariff as 'all customers' (described above).

Figure 5 shows the average top-up amount for each ACORN group.

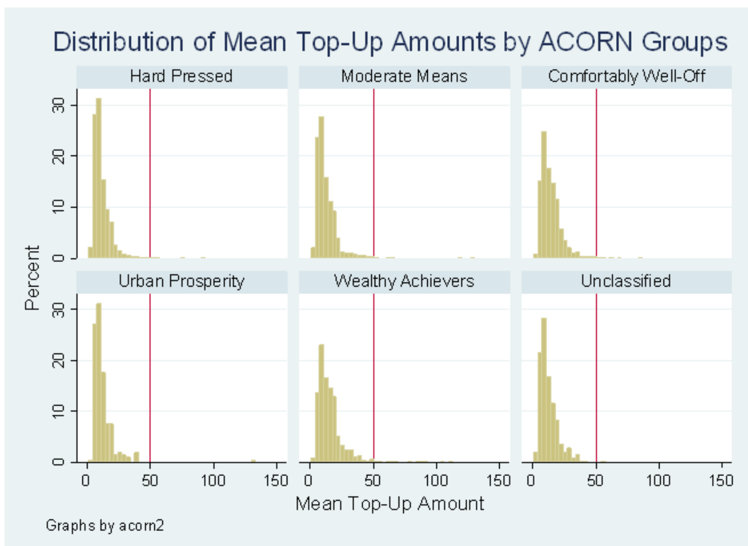

Figure 5: Distribution of Mean Top-up Amounts by ACORN Groups

Table 10 below shows the change in the number of top-ups and average top-up amount as a response to changes in tariff for ACORN groups 'comfortably well-off', 'urban prosperity', and 'wealthy achievers'.

	Change: tariff	Change: # of top-ups	Change: average top-up amount
06/08-07/08	0.01588	1218	-£0.08
	(14.0%)	(10%)	(-0.7%)
07/08-08/08	0	637	£0.21
	(0%)	(4.8%)	(1.9%)
08/08-09/08	0	-15	£0.31
	(0%)	(0%)	(2.7%)
09/08-10/08	0.04319	4111	£0.55
	(33.3%)	(29%)	(4.5%)
10/08-11/08	0	8	£0.62
	(0%)	(0%)	(5%)
11/08-12/08	0	-575	£1
	(0%)	(-3.1%)	(7.5%)
12/08-01/09	-0.01863	-877	-£0.87
	(-10.8%)	(-4.8%)	(-6.3%)
01/09-02/09	0	-498	-£0.22
	(0%)	(-2.9%)	(-1.7%)
02/09-03/09	0	-619	-£0.31
	(0%)	(-3.7%)	(-2.4%)
09/09-10/09	-0.768	-434	£0.18
	(-5.0)	(-2.8%)	(-1.3%)
10/09-11/09	0	394	£0.51
	(0%)	(2.6%)	(4.1%)

Table 10: Overall response to changes in Tariffs for highest ACORN groups

It is possible that the ACORN labels do not adequately describe the financial situation in each group and that even individuals classified as 'wealthy achievers' are to a large extent poor. If this were true, however, we would expect to find significant intra-month differences in top-up behaviour. This is not the case.[30]

[30] More specifically, to the extent that people tend to receive their salaries at the end of the month and social security benefits every fortnight, in a world with severe liquidity constraints

4.5 Expectations

Another possible explanation for the small top-up amounts is that customers are trying to game the top-up system. For example, because a change in tariff becomes effective only after a customer has entered a certain code into her meter, and because this code is typically combined with her vend code, it is possible that households are trying to sustain large stocks of credit on their meter (by making many small top-ups) to so smoothen an increase in tariff by postponing the first top-up after the increase in tariff.

Table 11 below shows the changes in tariff between 1.6.2008 and 30.11.2009. In order to test the hypothesis of 'hording' we compare the average number of top-ups and top-up amounts before the largest increase in tariff (on 1.10.2008; its announcement, respectively) and afterwards. We do not find any evidence for significant hording behaviour among a large number of customers. Instead, while the average top-up amount remained almost the same, as discussed earlier, the number of top-ups increased to account for the higher tariff.

Period	Unit Rate
01/06/08 - 30/06/08	£0.11373
01/07/08 - 30/09/08	£0.12961
01/10/08 - 31/12/08	£0.17280
01/01/09 - 30/09/09	£0.15417
01/10/09 -	£0.14649

Table 11: Changes in tariff over time

We also compare the average top-up amount between 1st July 2008 and 1st October 2008 with that between 1st October 2008 and 1st January 2009. The idea is to test whether the small top-up amounts can be explained by expectations of falling electricity tariffs. Again, we find no evidence for strategic behaviour.[31]

even among individuals in the highest ACORN groups, we would expect to find significantly higher average top-up amounts in the last ten days of a month than in the ten days before the last ten days; in the first week than in the second week, respectively.

However, when we compare the total top-up volume (number of top-ups and average top-up amount) in the last 10 days of each month with that in the ten days before that for individuals in the highest ACORN groups ('comfortably well off', 'urban prosperity', 'wealthy achievers') and in the first week of the month with that in the second week (for the same individuals), we find very modest evidence for liquidity constraints.

The change in the number of top-up trips in the last ten days (first week of the month) and the 10 days before that (second week) is -2.5% (-0.4%). The corresponding change in the average top-up amount is +3.5% (-1.8%)

[31] In addition, expectations of falling tariffs cannot explain why people adjust to changes in

4.6 Aversion to lose top-up

Yet another possible explanation for the small top-ups is that people are afraid of losing their top-up – either by losing their top-up voucher before entering it into their meter or through a technical problem of their meter.

One advantage of the new keypad system over the old system is that every keypad voucher works on exactly one meter only. This makes theft pointless. In addition, should a customer lose a top-up voucher, he/she can always get a reprint free of charge. (The same voucher code cannot be entered more than once). This excludes the possibility of (rational) aversion to lose one's top-up voucher as a possible explanation for our anomalies.

Similarly, there is little reason to believe that fear of losing money that is stored on one's meter as a result of a technical defect explains the small top-ups we observe: First of all, there is no evidence from the side of NIE Energy or customer protection groups that problems of this sort are common. In addition, even if a problem should occur, it is relatively straightforward for NIE Energy to get a sense of the loss involved from their records.

4.7 Commitment Device

Finally, it is possible that people use the small top-up amounts as a commitment device to save electricity. The idea can be illustrated as follows: Suppose an individual has the following utility function with (β, δ)-preferences (see Laibson, 1997 and O'Donoghue and Rabin, 1999):

$$U_t = u_t + \beta\delta u_{t+1} + \beta\delta^2 u_{t+2} + \beta\delta^3 u_{t+3} + ... \quad (3)$$

The main feature of these preferences is the β term which ensures – for β <1 – that the discounting between the present and the future is higher than between any future time periods.

The main implication of this model of individual preferences is that they introduce a conflict between an impatient 'present self' and a patient 'future self'. To give an example: Suppose electricity preservation has an immediate cost c at time 1 and a delayed pay-off b at t=2. How much does a customer want to preserve? The customer's preservation decision at t=0 – i.e. one period ahead – can be stated as follows: He/She will preserve if $\beta\delta^2 b_2 - \beta\delta c_1 \geq 0$; or $\delta^2 b_2 - \delta c_1 \geq 0$

tariffs almost exclusively by changing the number of top-ups.

Now, how much does a costumer actually preserve? At t=1 the customer preserves if $\beta\delta b_2 - c_1 \geq 0$ - which is smaller than the difference at t=0. So, compared to the desired, optimal level of preservation at t=0, a (β, δ)-customer preserves too little at t=1. The problem is that, because of the higher impatience of the 'present self' optimal contingent plans are not followed through.[32]

In response to the conflict between an impatient 'present self' and patient 'future self', a sophisticated agent will look for a commitment device. A commitment device is "an arrangement entered into by an agent which restricts his or her future choice set by making certain choices more expensive [...]" (Bryan et al, 2009). Commitment devices can take on a variety of forms. One form which might help us explain the small top-up amounts we observe is rationing (Wertenbroch, 1998). The idea is that customers self-impose a constraint on their consumption by rationing their purchasing quantities.

The rationing rule could say, for example: "Never buy more than £10 worth of top-up at a time and never buy top-up more than once a week". Such a rule implies that consumption at higher rates can only occur at the expense of incurring a cost, which is either the psychological cost of breaking one's own rule (see Thaler, 1985 or Heath and Loewenstein, 1991) or the cost of being without electricity until the next scheduled top-up.

It can be shown relatively easily that, if we assume that the psychological cost of breaking one's own rule is constant, then we can get a situation in which there is an upper bound on how long a top-up interval can be and so how large top-ups can be. The idea is that there are limits as to how strict a customer can make his or her commitment device by increasing the threat of being without electricity: if a payment schedule is long and the threat of being without electricity (at the end) high, a customer can always opt for breaking her own rule and so circumvent the higher cost of being without electricity.

The problem with this argument is that, while it can explain why customers choose relatively short top-up intervals and, by extension, small top-up amounts, it is inconsistent with people's adjustment behaviour to changes in tariff: The commitment argument implies an optimal top-up amount for a given tariff. The way people adjust to changes in tariff, however, implies that they choose different top-up amounts at a given tariff depending on the timing. To give an example,

[32] This line of reasoning has been used to explain seemingly irrational behaviour in a broad range of circumstances, ranging from saving too little for retirement, eating too much chocolate and not going to the gym often enough (Harris and Laibson, 2003; Wertenbroch, 2002; Dellavigna and Malmendier, 2004).

because of the asymmetric way people adjust to increases and decreases in tariff, a customer will choose a different top-up amount given tariff X at time A and B, if between A and B the tariff changes from X to Y and back to X.

5 Exploring the link between people's top-up behaviour and their electricity consumption

In the last sections we raised two empirical puzzles. We found that:

1. having a keypad meter tends to increase (rather than decrease) electricity consumption - despite information feedback on households' electricity use and higher transaction costs;

2. households tend to purchase very small top-up amounts every time they purchase top-up and adjust to increases in tariff almost exclusively by increasing the number of top-up trips

In this section, we explore whether there is a link between these two findings. Specifically, we ask whether and to what extent the effect of pre-payment metering on electricity consumption can be explained in the same way as people's top-up behaviour.

5.1 Cost Salience

One explanation for our two anomalies is that people perceive costs differently depending on whether they pay, say, 10 times £10 or £100 once: If paying 10 times £10 feels more trivial than paying £100 once, people might end up using more (rather than less) electricity under a pre-payment scheme - which allows them to disaggregate their electricity spending in whatever way they want.

Similarly, if paying 10 times £10 feels more trivial than paying £100 once, people (interested in minimising the negative hedonic impact of paying) can be expected to prefer relatively small top-ups to larger top-ups and to prefer adjusting to increases in tariff by increasing the number of top-up trips rather than the average top-up amount.

There is little systematic research on this idea[33] - with the exception of Gourville's work on the 'pennies-a-day-strategy'. In a series of experiments,

[33] The idea that the same amount can be perceived differently, depending on whether it is disaggregated or aggregated, features in several studies. Hardly ever, however, it is made explicit. Instead, in most cases (Gourville, 1998; Prelec and Loewenstein, 1998; Klee, 2006) it

Gourville (1998, 2003) showed that framing a donation request of £100 as 'mere 27p a day' is effective: he finds that the percentage of subjects agreeing to donate is significantly higher when they are asked to give up '27 p a day' compared to (the equivalent) £100 a year.

There is also some anecdotal evidence which suggests that people may perceive costs differently depending on the level of aggregation of the payment. Loewenstein and O'Donohue (2006) provide the following example:

"Consider two means of borrowing: (1) take out a loan of $10,000 on January 1 to be used for purchases over the next 12 months; or (2) slowly accumulate a credit card balance of $10,000. The standard economic model would say that a person ought to be (roughly) indifferent between the two options. But we suspect that many customers who accumulate $10,000 of credit card debt over a year would not have been willing to take out a $10,000 loan at the start of the year. Intuitively, they do not want to borrow and spend an extra $10,000, and so when faced with an aggregate decision of how much to borrow this year, they would choose much less. But when they make a series of disaggregated small borrowing decisions, people often end up borrowing a lot."

5.2 Linking Top-up Behaviour and Electricity Consumption

One testable implication of our hypothes is that there should be a link between people's top-up behaviour and their electricity consumption: if smaller top-up amounts are perceived as more trivial, we should find that an (exogenous) increase in top-up amount should lead to (an increase in cost salience and) a decrease in electricity use.

To assess the link between an (exogenous) increase in top-up amount and electricity use, we analyse the effect of an increase in the minimum top-up

is implicitly assumed that the main features of Kahneman and Tversky's (1979) value function – which are that people prefer to: i) separate gains (because the gain function is concave); ii) integrate losses (because the loss function is convex); integrate smaller losses with larger gains (to offset loss aversion); and segregate small gains (silver lining) from large losses (because the gain function is steepest at the origin, the utility of a small gain can exceed the utility of slightly reducing a large loss) – apply to general market transactions (such that people code the acquisition of a good as a gain and the forgone money as a loss).

Both Kahneman and Tversky (1984) and Thaler (1985, 1999) reject this idea. Thaler (1999) provides the following example of why he considers viewing costs as losses as descriptively inaccurate: "consider a thirsty consumer who would rather have a can of soda than one dollar and is standing in front of a vending machine that sells soda for 75 cents. Clearly the purchase makes her better off, but it might be rejected if the payment were cognitively multiplied by 2.25 (an estimate of the coefficient of loss aversion)".

amount. The change in minimum top-up took place on 15 May 2009. It applied only to top-ups purchased online or via a call centre and meant an increase in minimum top-up from £2 to £15.

5.3 Empirical framework

One simple regression method to evaluate the effect of a change in minimum top-up is based on a comparison of electricity consumption of those who had to change their normal top-up behaviour before and after the change in minimum top-up. For example, consider:

$$y_{it} = \alpha + \beta D_t + \varepsilon_{it}, i = i, ..., N; t = 0, 1... \quad (5)$$

Where D_t=1 in period 1 (post-intervention), D_t=0 in period 0 (pre-intervention), and y_{it} measures electricity consumption. The regression estimated from this model will yield an estimate of the policy impact parameter β. However, for this parameter to be consistent, we need that our sample remains comparable over time.

If we allowed α to vary between the two periods – e.g. as a result of seasonal variation – β would be confounded with this change and no longer reflect solely the effect of a change in minimum top-up on consumption. That is, looking at how consumption changed before and after a change in minimum top-up is uniformative if consumption would have changed even in the absence of the change in minimum top-up.

One way to improve on this design is to include an additional comparison group; one not impacted by the change in minimum top-up. The idea is that, if this comparison group is similar to the group which is affected by the change in minimum top-up (treatment group), it can tell us something about how electricity consumption of our treatment group would have evolved, had it not been affected by the change in minimum top-up.

Formally, using Meyer's (1995) notation, the relevant regression now is:

$$y_{it}^j = \alpha + \alpha_1 D_t + \alpha^1 D^j + \beta D_t^j + \varepsilon_{it}^j, i = 1, ..., N; t = 0, 1, ... \quad (6)$$

Where j is the group superscript, D_j=1 if j equals 1 and D_j=0 otherwise, D_{tj}=1 if both j and t equal 1 and D_{tj}=0 otherwise. ε is a zero-mean constant-variance error term. The equation does not include covariates, but they can be added.

What this relation implies is that, for the group for whom the change in minimum top-up was binding, we have pre-intervention

$$y_{i0}^1 = \alpha + \alpha^1 D^1 + \varepsilon_{i0}^1 \quad (7)$$

And post-intervention

$$y_{i1}^1 = \alpha + \alpha_1 + \alpha^1 D^1 + \beta + \varepsilon_{i1}^1 \quad (8)$$

Which suggests an impact of

$$y_{i1}^1 - y_{i0}^1 = \alpha_1 + \beta + \varepsilon_{i1}^1 - \varepsilon_{i0}^1 \quad (9)$$

The corresponding equations for the 'untreated' group (within our framework) are

$$y_{i0}^0 = \alpha + \varepsilon_{i0}^0 \quad (10)$$

And

$$y_{i1}^1 = \alpha + \alpha^1 + \varepsilon_{i0}^0 \quad (11)$$

Which suggests

$$y_{i1}^0 - y_{i0}^0 = \alpha_1 + \varepsilon_{i1}^0 - \varepsilon_{i0}^0 \quad (12)$$

The first difference shows the problem from before: rather than providing an estimate of the impact of a change in minimum top-up (β), what we get when comparing our 'treated' group before and after the change in minimum top-up is a composite of some other change (α_1) and the actual impact (β).

The second difference provides us with a way of getting rid of this bias: As the first difference it includes the period-1 specific effect α_1, which means that we can eliminate α_1 from the first difference by taking the difference between Equations (9) and (12):

$$(y_{i1}^1 - y_{i0}^1) - (y_{i1}^0 - y_{i0}^0) = \beta + (\varepsilon_{i1}^1 - \varepsilon_{i0}^1) - (\varepsilon_{i1}^0 - \varepsilon_{i0}^0) \quad (13)$$

Assuming that E[($\varepsilon_{i1}^1 - \varepsilon_{i0}^1$)-($\varepsilon_{i1}^0 - \varepsilon_{i0}^0$)] equals zero, we can obtain an unbiased estimate of β by the sample average of $(y_{i1}^1 - y_{i0}^1) - (y_{i1}^0 - y_{i0}^0)$. This method uses differences in differences. The identifying assumption of the approach is that the time trend in consumption of the treatment group and comparison group is the same in the absence of the change in minimum top-up.

5.4 Sample Description

Our discussion in the last section suggest that to assess the effect of a change in minimum top-up, we need to know for whom the change has lead to a change in top-up behaviour – i.e. we need to identify our 'treatment group'. In theory it is clear who this is: All those who would have purchased top-up online worth less than £15, but now purchase top-ups (still online) worth £15 or more.

The practical difficulty is that we do not observe what people would have done, had there not been a change in minimum top-up. However, because quite a large fraction of customers always top-up the same amount using the same payment channel, we can get a good sense of who would have purchased top-up online < £15 by looking at what people did in the three months before the change in minimum top-up.

In order to identify our 'treatment group', we divide our sample along the following dimensions: First, we separate out customers who tend to purchase top-up online. (This includes top-ups purchased via the internet and call centres). We say that a household purchases top-up online, if in the three months before the change in minimum top-up, it makes 50% or more of its purchases using the internet (or a call centre).

Secondly, we split the sample by median top-up amount – dropping all observations with a median top-up amount > £15 in the three months before the change in minimum top-up. Finally, in order to leave only people in the sample who tend to spend more or less the same amount every time they purchase top-up, we exclude all those for whom the mean absolute deviation from the mean is larger than 1.5 in the three months before the change in minimum top-up.[34] This leaves us with a sample size of 162. (See Table 12 for details).

Now that we have defined our 'treatment group', we need to find a group which is similar to our 'treatment group' except for the fact that it has not been affected by the change in minimum top-up. A natural candidate is 'Medium Top-uppers': We define it like our 'treatment group' with the exception that people in this group tend to make top-ups between £15 and £30 (rather than between £2 and £15). Because the change in minimum top-up has been from £2 to £15, Medium Top-uppers" are unlikely to have been affected by the change in minimum top-up. Table 12 below summarizes the two groups.[35]

[34] MAD of 1.5 is the minimum MAD before the sample size drops dramatically.

[35] We have also tried alternative comparison groups - such as 'Offline Top-uppers' - which we defined in the same way as our 'treatment group' except that people in this group tend to make their top-ups offline (i.e. using Paypoints, Payzones and/or Post offices) rather than

	Small top-ups: Median top-up <£15	Medium top-ups: Median top-up £15<£30
Online customers	Individuals: 162	Individuals: 298
	Mean amt: 10.69 (0.10)	Mean amt: 21.27 (0.11)
	Mean amt2: 9.87 (0.20)	Mean amt2: 21.85 (0.22)

Table 12: Sample description Treatment vs Comparison Group

The table shows the number of customers in each sub-sample, their mean top-up amount over the whole period (mean amt) and the mean top-up amount in the three months prior to the increase in minimum top-up (mean amt2).

5.5 Graphical Analysis

Because keypad meters do not require meter readings, we have no direct information on households' electricity consumption. What we do know is i) when households purchase top-up ii) how much top-up they purchase and iii) what their tariff is at any point in time. Assuming that households tend to purchase top-up always when they have roughly the same amount of credit left on their meter, this information allows us to work out what their consumption is.

In Figure 6, we show the average electricity consumption for all individuals calculated in this way.[36]

the internet. However, a simple logistic analysis - using data from the Northern Ireland Neighbourhood Information Service - suggets that while there are no significant differences between our 'treatment group' and 'Medium Top-uppers' there are considerable cross-sectional differences between our 'treatment group' and 'Offline Top-uppers'. In addition, a plot of pre-intervention consumption showed more similarlity between 'Medium Top-uppers' and our 'treatment group' than between 'Offline Top-uppers' and our 'treatment group'. The results from these analyses are available upon request.

[36] One possible difficulty that might arise from the construction of our data in this way is that the calculation of monthly consumption figures becomes increasingly imprecise as top-up amounts increase. The reason is that, ceteris paribus, the probability for a top-up to start and end in the same month decreases as top-up amounts increase which makes it harder to attribute top-up to a particular month. Given the small top-ups, we do not expect this to significantly alter our results. (Top-ups between £10 and £15 tend to take place on a weekly schedule).

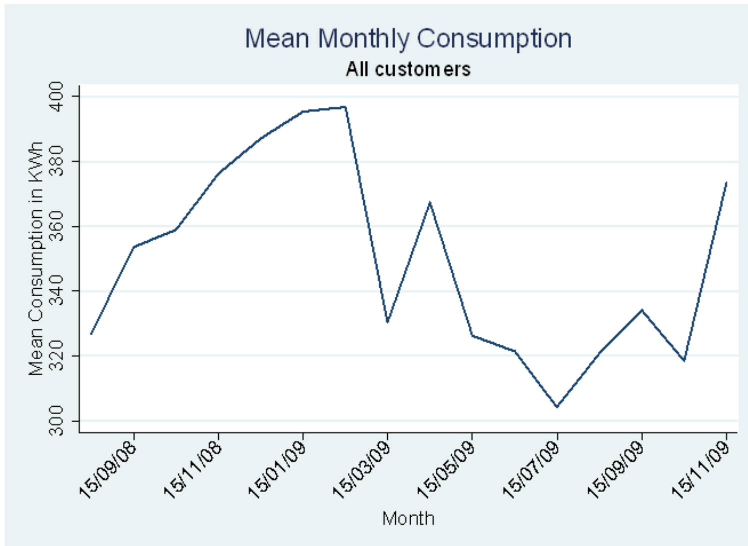

Figure 6: Mean Monthly Consumption

Figures 7 shows the average monthly electricity consumption for our main sub-samples: 'Treatment Group'; and 'Medium Top-uppers'. What we find is that consumption in the two groups follows each other relatively closely before the change in minimum top-up. After the introduction of the new minimum top-up amount, however, we find a decrease in consumption in the 'treatment group', while consumption for 'Medium Top-uppers' increases.

This finding is in line with our expectation: As people have to purchase larger amounts of top-up, they start paying more attention to their electricity spending and, hence, decrease their consumption.

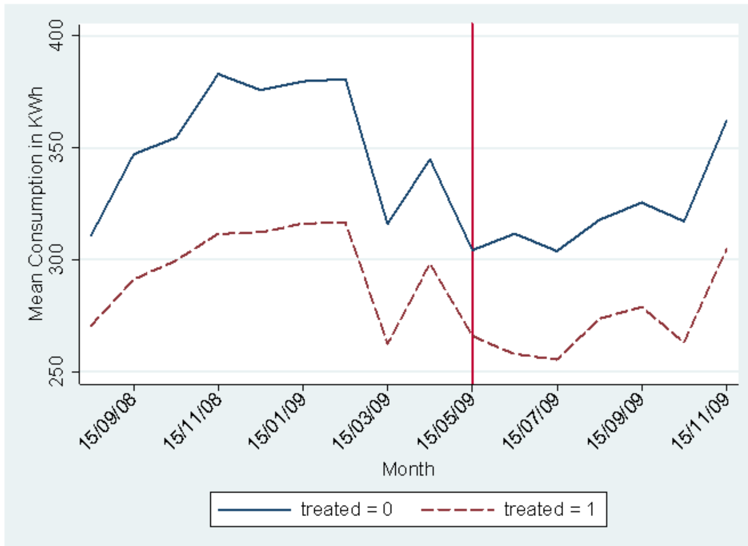

Figure 7: Mean Montly Consumption - Treatment and Comparison Group

One alternative explanation for our finding is that it is driven by 'liquidity constraints'. As discussed earlier, some people may be bad at saving for larger top-ups and so be forced to 'stretch' their credit over a longer period of time as a result of the change in minimum top-up. Yet, this possibility seems unlikely: Customers always have the possibility to just purchase their top-up offline (with a minimum top-up amount of £2 at the time).

In addition, looking at the distribution of customers in our 'Treatment Group', we find that customers in ACORN category 'hard pressed' are largely under-represented (28%), while customers in ACORN categories 'comfortably well-off' (25%) and 'wealthy achievers' are over-represented (28%).

In the next section, we go through the analysis again using a regression analysis. This allows us to determine how likely it is that our findings are due to chance only and allows us to test the appropriateness of the research design.

5.6 Regression Analysis

Let average monthly consumption be denoted by Y. Variables ST is equal to one if a customer typically purchases small top-ups (as defined earlier) – zero otherwise. TT is equal to 1 after the increase in minimum top-up – zero otherwise. Let X denote a vector of additional covariates. We estimate variants of the following linear model, which generalises the Difference-in-Difference estimation strategy set out earlier:

$$Y = \alpha + \beta_1 TT + \beta_2 ST + \gamma_1 TT * ST + \xi X + \varepsilon$$

In this specification, the second level interaction (γ_1) captures the treatment effect of the natural experiment. It equals the DD estimate. As a reference, specification 1 in Table 13 reports our treatment effect with no controls. Specification 2 replicates 1, controlling for top-up amount and month fixed effects. Specification 3 includes a set of control variables. For this, we match our data with information from the Northern Ireland Neighbourhood Information Service (NINIS).[37] All three specifications show that the change in minimum top-up led to a significant reduction in electricity consumption in our treatment group relative to the comparison group.

More specifically, what we find is that electricity consumption in our treatment group decreased by about 15 KWh as a result of the change in minimum top-up. Assuming the model is correctly specified the probability that such a result could have happened by chance is only ca 5%. In specification 4, we estimate an analogous model in logs instead of levels. An advantage of the log specification is that it may be a better model for comparison across groups with different baseline quantities. The log specification yields a slightly larger estimate than the levels models: a decline in monthly electricity consumption of 7%.

[37] NINIS is a census-based, locational data-set that contains information on a large number of socio-economic variables.

Variable	Spec (1)	Spec (2)	Spec (3)	Spec (4)	Spec (5)
Treatment	-10.84*	-14.66*	-15.55**	-0.07**	-19.13*
	(6.15)	(7.80)	(7.84)	(0.036)	(10.20)
Top-up amount FE		-37.11***	-41.34***	-0.17***	-44.06***
		(14.17)	(14.27)	(0.054)	(18.31)
Month FE		8.01*	8.65*	0.023**	13.79*
		(4.64)	(4.64)	(0.022)	(5.66)
Higher Prof. Occ.			1.06		
			(1.99)		
Lower Prof. Occ.			-0.17		
			(0.55)		
Unemployment rate			-1.63		
			(3.12)		
Over 60			0.26		
			(1.28)		
Married			1.07		
			(0.76)		
Children			1.19		
			(0.86)		
High Level of Edu.			-0.94		
			(2.12)		
No Education			1.25		
			(1.46)		

Table 13: Dif-in-Dif Estimates.
Statistically Significant at 1% ***; 5% **; 10% *.

To further probe into the possibility that our estimates are driven by liquidity constraint, we re-estimate our model using only individuals from high ACORN groups (comfortably well-off, urban prosperity, wealthy achievers). We find a statistically significant estimate in the ball-park of the estimates before (specification 5).

5.7 Robustness Checks

A concern in Difference-in-Difference analysis is that the model is mis-specified, insofar as the comparison group does not provide a good sense of what consumption would have been, had there not been an intervention (Bertrand, 2002). To check this possibility, we choose the month before and after the change in minimum top-up; then estimate (2) pretending that the chosen time period is the intervention period.

The idea is to test for a zero effect where it is known that the effect should be zero: finding a non-zero effect in these siutations would cast doubt on how suitable our comparison group is. Table 14 provides our estimation results. They show no significant change in electricity use in the month before and after the change in minimum top-up.

Building on the logic underlying this specification check, we estimate (2) using 'Treatment' and 'Medium Top-uppers' as before – with the exception that we now look at people purchasing their top-ups predominantly offline. As can be seen in Table 17 below, we do not find a significant change in electricity consumption for our pseudo 'Treatment Group'. This addresses the possibility that our findings for the actual 'Treatment Group' are driven by a shock at the time of the change in minimum top-up that differentially affected consumption of customers purchasing small top-ups.

Variable	Spec (Before)	Spec (After)	Spec (Offline)
Treatment	6.71	4.55	-1.85
	(8.67)	(8.10)	(2.94)
Top-up amount FE	-50.92***	-56.34***	-70.12***
	(14.77)	(14.54)	(5.32)
Month FE	-42.18***	-9.09 *	0.56
	(5.13)	(4.84)	(2.58)

Table 14: Placebo Estimates.

Statistically Significant at 1% ***; 5% **; 10% *.

5.8 Conclusion and Next Steps

In this paper we explored the role of pre-payment in the household energy consumption context. In the first part, we provided background to the metering situation in Northern Ireland. We discussed the large uptake of pre-payment metering after the introduction of keypad meters – and explored the differences in background characteristics of households with and without a keypad meter.

In the second part of the paper, we evaluated the effect of the keypad meter on electricity consumption. We found that having a keypad meter tends to increase (rather than decrease) electricity consumption. This is despite the fact that the keypad provides information feedback on electricity use and comes with higher transaction costs.

In the third part, we explored people's top-up behaviour (looking for a cue for our earlier finding): We noted that people tend to purchase relatively small top-up amounts, every time they purchase top-ups and adjust to increases in tariffs primarily by increasing the number of top-ups (rather than average top-up amounts).

In the final part of the paper, we suggested that both, the positve effect of using a keypad meter on electricity use and the puzzling top-up behaviour can be explained by the idea that people perceive costs differently depending on how aggregated they are. To test this idea, we explored the relationship between people's top-up behaviour and their electricity consumption. We exploited a change in minimum top-up to get an estimate of this relationship.

In line with our hypothesis, we found that an increase in the minimum top-up amount is associated with a decrease in electricity consumption. At least two policy relevant questions arise from our analysis:

- Should we discourage people from using pre-payment and encourage post-consumption payment, instead? or
- To the extent that there is a preference for pre-payment metering, should we encourage pre-payment customers to top-up larger amounts every time they purchase top-up?

Taking this work forward, the three main tasks will be: i) to try to get better data on electricity consumption; ii) to model our main argument more rigorously and iii) to test, in a large-scale field experiment, how the relationship between top-up behaviour and electricity consumption varies across different parts of the population; for different changes in top-up amount; and over longer and shorter periods of time.

6 Bibliography

References

[1] Babin, B. J., Darden, W. R., Griffin, M., & Darden, R. (2010). Work and / or Fun : Shopping Measuring Value Hedonic and Utilitarian, 20(4), 644-656.

[2] Baumol, W. J. (1952). The Transactions Demand for Cash: An Inventory Theoretic Approach. The Quarterly Journal of Economics, 66(4), 545 - 556. Retrieved from http://www.jstor.org/stable/1882104.

[3] Becker, G. (1965). A Theory of the Allocation of Time. The economic journal, 75(299), 493–517. JSTOR. Retrieved from http://www.jstor.org/stable/2228949.

[4] Bertrand, M., Duflo, E., & Mullainathan, S. (2004). How Much Should We Trust Differences-in-Differences Estimates? The Quarterly Journal of Economics, 119(1), 249 - 275. Retrieved from http://www.jstor.org/stable/25098683.

[5] Brophy Haney, A.; Pollitt, M.; Jamasb, T. (2009). Smart Metering and Electricity Demand: Technology, Economics and International Experienceitle. EPRG Working Paper Series, (EPRG0903).

[6] Bryan, G.; Karlan, D.; Nelson, S. (2009). Commitment Devices. Science, 1-37.

[7] Burke, R. R., Harlam, B. A., Kahn, B. E., & Lodish, L. M. (1992). Comparing Dynamic Consumer Choice in Real and Computer-Simulated Environments. The Journal of Consumer Research, 19(1), 71 - 82. Retrieved from http://www.jstor.org/stable/2489189.

[8] Caplin, A., & Schotter, A. (2008). The foundations of positive and normative economics: a handbook. Oxford University Press, USA.

[9] Chetty, R, Looney, A & Kroft, K (2009). Salience and taxation: Theory and evidence. The American Economic Review, 99(4), pp. 1145-1177.

[10] Darby, S. (2006). No TitleThe effectiveness of feedback on energy consumption. Envirnmental Change Institute Oxford.

[11] Dehejia, R. H., & Wahba, S. (1999). Causal Effects in Nonexperimental Studies: Reevaluating the Evaluation of Training Programs. Journal of the American Statistical Association, 94(448), 1053 - 1062. Retrieved from http://www.jstor.org/stable/2669919.

[12] Dehejia, R. H., & Wahba, S. (2002). Propensity Score-Matching Methods for Nonexperimental Causal Studies. The Review of Economics and Statistics, 84(1), 151 - 161. Retrieved from http://www.jstor.org/stable/3211745.

[13] Dellavigna, S., & Malmendier, U. (2004). Contract Design And Self-Control: Theory And Evidence*. Quarterly Journal of Economics, 119(2), 353-402.

[14] Dulleck, U. (2004). Do customer information programs reduce household electricity demand?—the Irish program. Energy Policy, 32(8), 1025-1032. doi: 10.1016/S0301-4215(03)00060-0.

[15] Feinberg, R. A. (1986). Credit Cards as Spending Facilitating Stimuli: A Conditioning Interpretation. The Journal of Consumer Research, 13(3), 348 - 356. Retrieved from http://www.jstor.org/stable/2489426.

[16] Finkelstein, A. (2007). E-Ztax: Tax Salience and Tax Rates. SSRN eLibrary. SSRN. doi: 10.2139/ssrn.964887.

[17] Gourville, J. (1998). Pennies-a-day: The effect of temporal reframing on transaction evaluation. Journal of Consumer Research. Retrieved from http://www.journals.uchicago.edu/doi/abs/10.1086/209517.

[18] Gourville, J. T. (2003). The Effects of Monetary Magnitude and Level of Aggregation on the Temporal Framing of Price. Marketing Letters, 14(2), 125-135. doi: 10.1023/A:1025467002310.

[19] Gourville, J. T., & Soman, D. (1998). Payment Depreciation: The Behavioral Effects of Temporally Separating Payments from Consumption. The Journal of Consumer Research, 25(2), 160 - 174. Retrieved from http://www.jstor.org/stable/2489598.

[20] Halvorsen, B. (2001). The flexibility of household electricity demand over time. Resource and Energy Economics, 23(1), 1-18. doi: 10.1016/S0928-7655(00)00035-X.

[21] Harris, Ch; Laibson, D. (2003). Hyperbolic Discounting and Consumption. Advances in economics and econometrics: theory and applications, Eighth Wor.

[22] Heath, C. (1995). Escalation and De-escalation of Commitment in Response to Sunk Costs: The Role of Budgeting in Mental Accounting. Organizational Behavior and Human Decision Processes, 62(1), 38-54. doi: 10.1006/obhd.1995.1029.

[23] Heckman, J. J., Ichimura, H., & Todd, P. (1998). Matching as an Econometric Evaluation Estimator. The Review of Economic Studies, 65(2), 261 - 294. Retrieved from http://www.jstor.org/stable/2566973.

[24] Hirschman, E. C. (1979). Differences in Consumer Purchase Behavior by Credit Card Payment System. The Journal of Consumer Research, 6(1), 58 - 66. Retrieved from http://www.jstor.org/stable/2488726.

[25] Ichino, A, Mealli, F, Nannicini, T. (2008): From Temporary Help Jobs to Permanent Employment: What Can We Learn from Matching Estimators and their Sensitivity? Journal of Applied Econometrics, 23, 305–327.

[26] Kahneman, D., & Tversky, A. (1979). Prospect Theory: An Analysis of Decision under Risk. Econometrica, 47(2), 263 - 292. Retrieved from http://www.jstor.org/stable/1914185.

[27] LIJESEN, M. (2007). The real-time price elasticity of electricity. Energy Economics, 29(2), 249-258. doi: 10.1016/j.eneco.2006.08.008.

[28] Laibson, D. (n.d.). Golden Eggs and Hyperbolic Discounting. The Quarterly Journal of Economics, 112(2). Retrieved from http://www.jstor.org/stable/2951242.

[29] Laibson, D. (1997). Golden Eggs and Hyperbolic Discounting*. Quarterly Journal of Economics, 112(2), 443-477. doi: 10.1162/003355397555253.

[30] Loewenstein, G., & Prelec, D. (1992). Anomalies in intertemporal choice: Evidence and an interpretation. The Quarterly Journal of Economics, 3(4), 573-597.

[31] Mankiw, N., & Summers, L. (1986). Money demand and the effects of fiscal policies. Journal of Money, Credit and Banking. Retrieved from http://www.jstor.org/stable/1992462.

[32] Morwitz, V. G., Greenleaf, E. A., & Johnson, E. J. (1998). Divide and Prosper: Consumers' Reactions to Partitioned Prices. Journal of Marketing Research, 35(4), 453 - 463. Retrieved from http://www.jstor.org/stable/3152164.

[33] O'Donoghue, T., & Rabin, M. (1999). Doing It Now or Later. The American Economic Review, 89(1), 103 - 124. Retrieved from http://www.jstor.org/stable/116981.

[34] Owen, G., & Ward, J. (2010). Smart pre-payment in Great Britain. Sustainability First. Review Literature And Arts Of The Americas, (March), 1-36.

[35] Owen, G., Ward, J., Hobson, J., Osborn, D., Pugh, T., Sigsworth, D., et al. (2007). Smart meters in Great Britain : the next steps ? Energy, (July).

[36] O'donoghue, T., & Rabin, M. (1999). Doing it now or later. American Economic Review, 114(3), 103-124.

[37] Prelec, D., & Loewenstein, G. (1998). The Red and the Black: Mental Accounting of Savings and Debt. Marketing Science, 17(1), 4 - 28. Retrieved from http://www.jstor.org/stable/193194.

[38] Prelec, D., & Loewenstein, G. (1998). The red and the black: Mental accounting of savings and debt. Marketing Science, 17(1), 4–28. JSTOR. Retrieved from http://www.jstor.org/stable/193194.

[39] Romer, D. (1986). A Simple General Equilibrium Version of the Baumol-Tobin Model. The Quarterly Journal of Economics, 101(4), 663 - 686. Retrieved from http://www.jstor.org/stable/1884173.

[40] Rosenbaum, P. R., & Rubin, D. B. (1983). The Central Role of the Propensity Score in Observational Studies for Causal Effects. Biometrika, 70(1), 41 - 55. Retrieved from http://www.jstor.org/stable/2335942.

[41] Smith, J. A., & Todd, P. E. (n.d.). Reconciling Conflicting Evidence on the Performance of Propensity-Score Matching Methods. The American Economic Review, 91(2). Retrieved from http://www.jstor.org/stable/2677743.

[42] Soman, D. (1998). The Illusion of Delayed Incentives: Evaluating Future Effort-Money Transactions. Journal of Marketing Research, 35(4), 427 - 437. Retrieved from http://www.jstor.org/stable/3152162.

[43] Soman, D. (2003). The Effect of Payment Transparency on Consumption: Quasi-Experiments from the Field. Marketing Letters, 14(3), 173-183. doi: 10.1023/A:1027444717586.

[44] Soman, D., & Cheema, A. (2002). The Effect of Credit on Spending Decisions: The Role of the Credit Limit and Credibility. Marketing Science, 21(1), 32 - 53. Retrieved from http://www.jstor.org/stable/1558056.

[45] Soman, D., & Gourville, J. T. (2001). Transaction Decoupling: How Price Bundling Affects the Decision to Consume. Journal of Marketing Research, 38(1), 30 - 44. Retrieved from http://www.jstor.org/stable/1558569.

[46] Soman, D., & Lam, V. M. (2002). The Effects of Prior Spending on Future Spending Decisions: The Role of Acquisition Liabilities and Payments. Marketing Letters, 13(4), 359-372. doi: 10.1023/A:1020374617547.

[47] Soman, D; Cheema, A. (2008). The Effect of Partitions on Controlling Consumption. Journal of Marketing Research, 45(6), 665-675. doi: 10.1509/jmkr.45.6.665.

[48] Thaler, R. (1985). Mental Accounting and Consumer Choice. Marketing Science, 4(3), 199 - 214. Retrieved from http://www.jstor.org/stable/183904.

[49] Thaler, R. H. (1991). ‘Some Empirical Evidence on Dynamic Inconsistency. Quasi rational economics, 1, 127-136.

[50] Thaler, R. H. (2000). From Homo Economicus to Homo Sapiens. The Journal of Economic Perspectives, 14(1), 133 - 141. Retrieved from http://www.jstor.org/stable/2647056.

[51] Trope, Y., & Fishbach, A. (2000). Counteractive self-control in overcoming temptation. Journal of Personality and Social Psychology, 79(4), 493-506.

[52] Wertenbroch, K. (1998). Consumption Self-Control by Rationing Purchase Quantities of Virtue and Vice. Marketing Science, 17(4), 317 - 337. Retrieved from http://www.jstor.org/stable/193203.

[53] Wertenbroch, K. (1998). Consumption self-control by rationing purchase quantities of virtue and vice. Marketing Science, 317-337.

[54] Wood, G. (2003). Dynamic energy-consumption indicators for domestic appliances: environment, behaviour and design. Energy and Buildings, 35(8), 821-841. doi: 10.1016/S0378-7788(02)00241-4.

[55] Zhang, T.; Nuttall, W. (n.d.). Evaluating Government's Policies on Promoting Smart Metering Diffusion via Agent Based Simulation. Journal of Product Innovation Management.