View metadata, citation and similar papers at core.ac.uk
brought to you by CORE
provided by Apollo

Self-Disconnection Among Pre-Payment Customers – A Behavioural Analysis

Philipp-Bastian Brutscher

March 2012

CWPE 1214

UNIVERSITY OF CAMBRIDGE
Electricity Policy Research Group

Self-Disconnection Among Pre-Payment Customers - A Behavioural Analysis

EPRG Working Paper 1207

Cambridge Working Paper in Economics 1214

Philipp-Bastian Brutscher

Abstract

In this paper, we revisit the problem of self-disconnection among pre-payment energy customers. Using metering data from 2.3 million electricity pre-payment customers, we study how often households with an electricity pre-payment meter tend to self-disconnect over the course of a year - and why they do so.

What we find is that, in any given year, the majority of households (ca. 78%) do not self-disconnect; ca. 12% self-disconnect once; ca. 3% self-disconnect more often than four times. We also find that most self-disconnections (ca. 62%) last for less than one day; between 72% and 82% last for less than two days; 12%-18% last for more than 3 days.

As for the main driver of self-disconnection, we identify financial constraints. This suggests that it is likely to be difficult/expensive to reduce the total number of self-disconnections. In the last part of the paper, we argue, however, that it may (still) be possible to reduce the negative impact of self-disconnection in a relatively inexpensive way - at least to some extent - by helping households to better smooth their self-disconnections over the course of a year.

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

EPRG WORKING PAPER

UNIVERSITY OF CAMBRIDGE

Electricity Policy Research Group

Keywords Pre-payment; Self-Disconnection: Commitment Device; Self Control; Fuel Poverty

JEL Classification D03, D04, D12, D14

EPRG WORKING PAPER

Contact pbb24@cam.ac.uk
Publication March 2012
Financial Support EPSRC Flexnet

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

EPRG No.1207

Self-Disconnection Among Pre-Payment Customers - A Behavioural Analysis.

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

March 07 2012

Abstract

In this paper, we revisit the problem of self-disconnection among pre-payment energy customers. Using metering data from 2.3 million electricity pre-payment customers, we study how often households with an electricity pre-payment meter tend to self-disconnect over the course of a year – and why they do so.

What we find is that, in any given year, the majority of households (ca. 78%) do not self-disconnect; ca. 12% self-disconnect once; ca. 3% self-disconnect more often than four times. We also find that most self-disconnections (>62%) last for less than one day; between 72% and 82% last for less than two days; 12%-18% last for more than 3 days.

As for the main driver of self-disconnection, we identify 'financial constraints'. This suggests that it is likely to be difficult/expensive to reduce the total number of self-disconnections. In the last part of the paper, we argue, however, that it may (still) be possible to reduce the negative impact of self-disconnection in a relatively inexpensive way – at least to some extent – by helping households to better smooth their self-disconnections over the course of a year.

*I would like to thank my supervisor, David Newbery, for reading and discussing numerous drafts of this paper. I would also like to thank Michael Pollitt for commenting on and helping to improve the paper at various stages of its development. Further, I would like to thank Richard Duck; Richard Jones; Christopher Mitchell; and Christoniel Putter from Centrica for their interest in the project and their great support. Last but not least, I am grateful to Julian Misell and Trevor Taylor from Ipsos Mori for their excellent collaboration on the field work. Financial support from the EPSRC (Flexnet) is gratefully acknowledged. The usual disclaimer applies.

1 Introduction

Pre-payment energy meters are widely used in Great Britain. To date approximately 3.6 million households (14%) use pre-payment meters to pay for their electricity; 2.6 million households (10%) use them to pay for their gas consumption; and 2.3 million households (9%) use them to pay for both: electricity and gas.

Modern pre-payment meters work in a similar way to pre-payment mobile phones: customers purchase credit at an outlet which they then use to 'top-up' their meters. The meters come with several advantages: they allow customers to break their energy bills into a series of (arbitrarily) small payments; they often come with a lower-than-standard tariff; and detailed information feedback allows pre-payment customers to better monitor their energy use.

The main disadvantage of pre-payment meters is that they can lead to 'self-disconnection': if credit runs out, energy supply is stopped.[1] Little is known about self-disconnection: while the regulator (OFGEM) records and monitors the number of disconnections among non-pre-payment customers, no comparable effort is made to capture the extent (and drivers) of self-disconnection among pre-payment customers.

Several studies have tried to estimate the extent and drivers of self-disconnection. The conclusions from these studies, however, vary significantly: while some find that self-disconnection is a widespread phenomenon, others argue that it is an issue for at most a small fraction of pre-payment customers. Similarly, while some studies argue that 'financial constraints' are at the core of most self-disconnections, others claim that 'financial constraints' are responsible for a relatively small share of self-disconnections only.

In this paper, we take a fresh look at the issue of self-disconnection. We ask:

- What is the extent of self-disconnection?
- What are the main drivers of self-disconnection? And
- What policy measures are there to address the issue of self-disconnection?

We add to the existing literature in several ways: we use (new) metering data from 2.3 million pre-payment customers to study how often households self-disconnect; we provide a series of formal hypothesis tests for why households

[1] We will provide a more precise definition of what we mean by self-disconnection in section 3.3.

self-disconnect (instead of relying on descriptive evidence only); and suggest several novel policy instruments which can help reduce the negative impact of self-disconnections.

The main findings of our research are that every year only about a fifth of pre-payment customers self-disconnect; many of these customers self-disconnect for a relatively short period of time only; for some of them, however, the (total) duration without electricity (due to self-disconnection) is considerable. As for the main driver of self-disconnection, we identify financial constraints.

From a policy perspective the main conclusion arising from our research is that although reducing the total number of self-disconnections is likely to be difficult/expensive, it may (still) be possible to reduce the negative impact of self-disconnection in a relatively inexpensive way by helping households to better smooth the number of self-disconnections over the course of a year (instead of letting them self-disconnect when it hurts the most).

The paper is organised as follows: In the first part, we give some background information on the pre-payment situation in Great Britain. The second part provides new estimates of the extent of self-disconnection in Great Britain. In the third part, we argue that reducing the total number of self-disconnections is likely to be difficult/expensive. The fourth, fifth and sixth part offer an alternative angle on the problem (by looking at households' timing of self-disconnections). The seventh part concludes.

2 Background

In this section, we provide some background information on the pre-payment situation in Great Britain.

2.1 Pre-Payment Metering in Great Britain

Pre-payment meters have experienced a significant increase in popularity in Great Britain in the past 20 years. Figure 1a illustrates the change in the use of pre-payment meters since 1990.

The figure shows an increase in the share of households using a pre-payment meter to pay for their electricity from ca 8% in 1990 to 14% in 2008 ; an increase in the share of households using a pre-payment meter to pay for their gas from 3% to 10%; and an increase in the share of households using a pre-payment

meter to pay for both electricity and gas from 2% to 9%.[2]

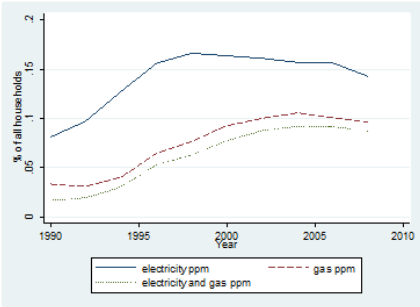

Figure 1a: Share of households using pre-payment metering.[3]

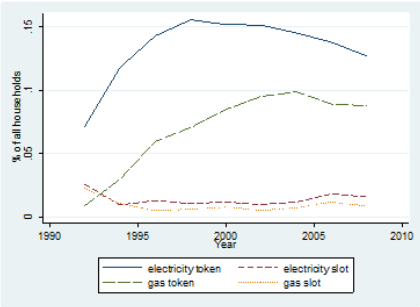

Figure 1b: Share of households by type of meter.

The recent increase in popularity of pre-payment meters is closely linked to a change in technology:[4] Figure 1b shows the evolution of the share of households using traditional coin-in-the-slot meters vs modern token/key meters. The figure shows that while the share of households using slot meters remained relatively constant, the share of households using token/key meters has increased significantly (since their introduction approximately 20 years ago).

The main difference between coin-in-the-slot meters and token/key meters is that with token/key meters money is not entered into the meters directly – but via pre-programmed cards (token meters); or a customer specific chip card (key meters), respectively. This has made/makes pre-payment metering safer with token/key meters – and so more acceptable for customers.

The cashless top-up mechanism of token/key meters also meant/means that no money has to be collected from the meters. This has made/makes pre-payment metering (with token/key meters) economically more attractive for energy companies – which, in turn, has allowed them to reduce their energy tariffs for these meters.

A second difference between coin-in-the-slot meters and token/key meters is that token/key meters typically come with an emergency credit facility: while households with a traditional coin-in-the-slot meter are without energy as soon

[2] Data: 2009 Living Cost and Food Survey (LCF). All statistics are weighted.

[3] Data: Family Expenditure Survey; Expenditure and Food Survey and Living Cost and Food Survey.

[4] For a detailed analysis of the diffusion process underlying pre-payment metering see Zhang et al (2009)

as they run out of credit, households with a token/key meter are typically provided with (on average) £5 of 'emergency credit'.

The idea of 'emergency credit' is to give households time to make another top-up before switching off their energy supply. Naturally, whatever emergency credit has been used will be subtracted from the next 'top-up(s)' – but no interest/penalty will be charged.

2.2 Who uses pre-payment meters?

Despite the increase in popularity of pre-payment meters in the past 20 years, it is important to note that pre-payment customers (still) represent a fairly distinct group of households.

Using data from the Living Cost and Food Survey 2009, Table 1 below reports summary statistics for households with no pre-payment meter; households with an electricity pre-payment meter (but no gas pre-payment meter); and for households with both: an electricity pre-payment meter and a gas pre-payment meter.[5]

[5] Data: 2009 LCF. All statistics are weighted.

Variable	Non-PPM	E-PPM	p-value (Non vs E-PPM)	E&G-PPM	p-value (E-vs E&G-PPM)
N	4,605	281		447	
Age	53.16	46.46	<0.01	44.24	0.04
	(16.81)	(15.54)		(15.65)	
Gender	0.36	0.45	<0.01	0.55	0.01
(Fem=1)	(0.48)	(0.50)		(0.50)	
Single	0.33	0.49	<0.01	0.49	0.72
(Sing=1)	(0.47)	(0.50)		(0.50)	
Income	704.93	443.22	<0.01	378.42	0.02
(weekly)	(475.76)	(270.86)		(236.75)	
Active	0.63	0.54	<0.01	0.43	0.01
(Act=1)	(0.48)	(0.50)		(0.50)	
Hh size	2.32	2.32	0.81	2.58	0.01
	(1.24)	(1.41)		(1.43)	
# Children	0.33	0.47	<0.01	0.62	0.03
	(0.73)	(0.88)		(0.95)	
Type house	0.27	0.05	<0.01	0.01	<0.01
(det=1)	(0.45)	(0.21)		(0.01)	
Rent	0.21	0.67	<0.01	0.77	<0.01
(Rent=1)	(0.41)	(0.47)		(0.42)	
Countrys.	0.14	0.15	0.83	0.01	<0.01
(Cntrys=1)	(0.35)	(0.36)		(0.11)	

Table 1: Summary Statistics by payment type.

The table shows that all three groups of households are (statistically) fairly different from each other. To give an example: the table shows that households with an electricity pre-payment meter are significantly more likely to have a female household head; significantly less likely to be economically active; and significantly more likely to rent their home than households without a pre-payment meter.

Similarly for households with an electricity and gas pre-payment meter: the table shows that households with an electricity pre-payment and gas pre-payment meter are (even) more likely to have a female household head; (even) less likely to be economically active; and (even) more likely to rent their home

than households with an electricity meter only.

3 The Extent of Self-Disconnection

In this section, we review the literature on self-disconnection – and provide new estimates of the extent of the problem.

3.1 Self-disconnection – What we know

Several studies have looked at the extent and drivers of self-disconnection. The main studies are summarized in Table 2 below.

Study	Appr.	Sample	Extent of SD	Drivers
Drakeford (1995)	Survey	388 households-recruited through the Citizen Adv. Bureaux and social services offices in the South of Wales.	51% (60%) of electricity (gas) customers report having been w/o supply in past	Primarily problems to afford credit.
Centre for Sustain. Energy (1998)	Survey	271 households - which receive council tax benefits	28% of electr./ gas report having self-disconnctd in the past.	62% self-disconnected because of a lack of money; 30% because of a closed outlet.
Doble (2000)	Face to Face interviews	200 randomly chosen pre-payment gas customers in Coventry area	33% of cust. report having self-disconnctd at one point in the last year.	Reasons for self-disconnections were primarily 'convenience'
Centre f. manage.. u. Regul. (2001)	Survey	100 households with electricity/gas pre-pay ment meter	14% of gas pre-payment; 9% of electricity prepayment meter customers	86% (91%) of electricity (gas) customers report self-disconnctn for a reason other than lack of money.
Regional Studies	Survey[6]	Between 48 and 2360 households in certain small regions.	Between 15% & a large number of pre-payment households' report self-disconnections.	Reasons vary from having forgotten to top-up meters to coordination problems to lack of money.

Table 2: Overview of the Self-Disconnection Literature

[6] Scottish Hydro-Electric (1997); Scottish Hydro-Electric (1998); Merseyside Right to Fuel Action and Liverpool City Council (1998); Birminham Settlement, Community Energy Re-

What is interesting about the previous literature is the variation in findings across studies. To give an example: while Drakeford (1995) finds that more than 50% of households self-disconnect at one point over the course of a year, Doble (2000) argues that this number is closer to 30%. Several regional studies come to even lower estimates.[7] The same is true with respect to the drivers of self-disconnection.[8]

It is not clear what is at the core of these differences across studies. It seems likely, however, that differences in the timing of studies play an important role (with more households self-disconnecting in the past few years than today). In addition – because all previous studies rely exclusively on survey data – it seems likely that differences in sample selection; sample size; survey method; and the context of the survey are responsible for (at least part of) the variation in findings across studies.

3.2 Metering Data

One way to deal with some of the methodological issues of the previous literature is by using metering data.

We have access to metering data from 2.3 million households with an electricity pre-payment account with British Gas – spanning the years 2007 to 2010 (inclusive). The dataset – which was provided to us in confidence – includes information on electricity consumption, top-up behaviour and the use of emergency credit.

The main advantage of metering data (vis-a-vis the survey data used in most earlier studies) is that it allows us to get access to a significantly larger sample of households than was possible in most earlier studies. In addition, metering data allows us to circumvent problems of recall/response bias; and it allows us to ensure that our findings are independent of the survey method used/the context of the survey.[9]

To get a sense for whether the households in our metering dataset are also (more) similar to a repesentative sample of pre-payment customers, we randomly

search and Bristol Energy Centre (1993)

[7] These include: Birmingham Settlement (1993); Liverpool City Council (1998); Money Matters and Scottish Power (1997); Scottish Hydro Electric (1998).

[8] While the study by the Centre for Sustainable Energy (1998) finds that 'lack of money' is the main driver of self-disconnection, a comparable study by the Centre for Management under Regulation (2001) argues that forgetting to purchase credit is of at least equal importance to understand why households self-disconnect.

[9] For the importance of survey method/context see e.g. Browning et al (2003)

drew a sample of households from our dataset; asked them a series of background questions (in a telephone survey); and compared the findings from this survey with the findings from the 2009 Living Cost and Food Survey.

Table 3 shows the results from this analysis.[10]

Variable	BG Sample	LCF Sample	p-value
Number of observations	502	728	
Age	45.02	45.07	0.95
	(13.73)	(15.64)	
Gender	0.63	0.51	<0.01
(Female=1)	(0.48)	(0.50)	
Marital status	0.37	0.49	<0.01
(Single=1)	(0.48)	(0.50)	
Income	190.60	209.021	0.22
	(200.14)	(291.7098)	
Laour Market Status	0.48	0.47	0.34
(Active=1)	(0.50)	(0.50)	
Household size	2.61	2.47	0.15
	(1.48)	(1.43)	
Number of Children in hh.	0.63	0.56	0.25
	(1.20)	(0.93)	
Type house	0.07	0.04	0.11
(detached=1)	(0.26)	(0.14)	
Rent	0.75	0.73	0.43
(Rent=1)	(0.43)	(0.44)	

Table 3: Summary Statistics British Gas Sample and LCF Sample

The Table shows that – except for gender and marital status[11] – the households in our sample are (statistically) very similar to the households with an electricity pre-payment meter in the Living Cost and Food Survey which suggests that – despite coming from one company only – our dataset is (close to being) representative for all households with an electricity pre-payment meter.

[10] All statistics are weighted – using survey weights.

[11] Please note: since we are targeting the 'household member in charge of purchasing top-up' (rather than the household reference person) it is not surprising that this person's background characteristics are different from those of the one interviewed in the LCF (i.e. the household reference person in the LCF).

3.3 Defining Self-Disconnection

One short-coming of our metering data is that it does not record directly when households are without electricity.

What is recorded, however, is whether (and how much) emergency credit a household has used by the time it tops-up its meter. This allows us to get a sense for whether and when households self-disconnect – defining a self-disconnection as a situation in which a household has used up all its emergency credit by the time it tops-up its meter.

Most earlier studies define self-disconnection as a situation in which energy supply is stopped 'due to insufficient funds on the meter'. This is less strict than our definition. The reason is that energy supply can stop (due to insufficient funds) before a household has used up its 'emergency credit'. To use emergency credit, it has to be activated (which can be done by pressing a key on the meter).

To the extent, however, that what matters from a policy perspective is when households are without the possibility of electricity (unless they purchase more credit), it seems reasonable to define self-disconnection as a situation in which supply is stopped and a household has exhausted all the credit available to it (i.e. including its 'emergency credit') – rather than one where electricity supply is stopped but emergency credit is possibly still available.

3.4 The Extent of Self-Disconnection Revisited

Figure 2 below shows how often households tend to self-disconnect over the course of a year – using our definition of self-disconnection.

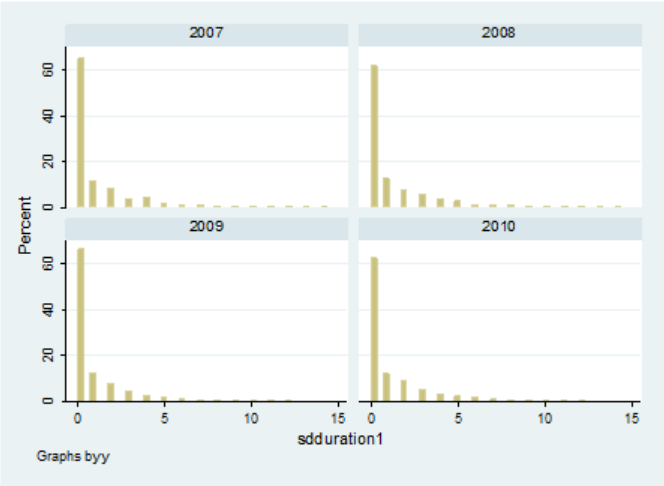

Figure 2: Distribution of the Frequency to Self-Disconnect.

The figure shows that in 2010 the large majority of households never self-disconnected (ca 78%); 12% of households self-disconnected once; approximately 3% of households self-disconnected more often than four times. The figure also shows that the pattern in self-disconnections is relatively stable across years – with a slightly lower share of households self-disconnecting in 2010 than the years before.

3.5 The Extent of Self-Disconnection Revisited 2

We can also use our metering data to proxy the duration of self-disconnections. Figure 3 below shows the distribution of when households top-up their meters after (being predicted to)[12] having self-disconnected.

[12] We have no direct information for how long households self-disconnect. However, we can proxy the duration of self-disconnections by taking the difference between how long the last top-up before a self-disconnection should have lasted and the date of the next 'top-up'. For this calculation we proceed as follows: first, we take the ratio between the consumption of each household that has self-disconnected and the median household (excluding households from the calcuation for the median which have self-disconnected) one month before the self-disconnection. We then assume that our households would have consumed the same fraction of what the median household consumes in the month in which our households actually self-disconnect (as in the month before). In the final step, we take difference between the prediction of when our households should have run out of credit resulting from the last step and when they actually top-up their meters - to proxy the duration of each self-disconnection. Please note: this is likely to be a relatively conservative measure of the duration of self-disconnections, since it does not take into account that households may be rationing before they self-disconnect.

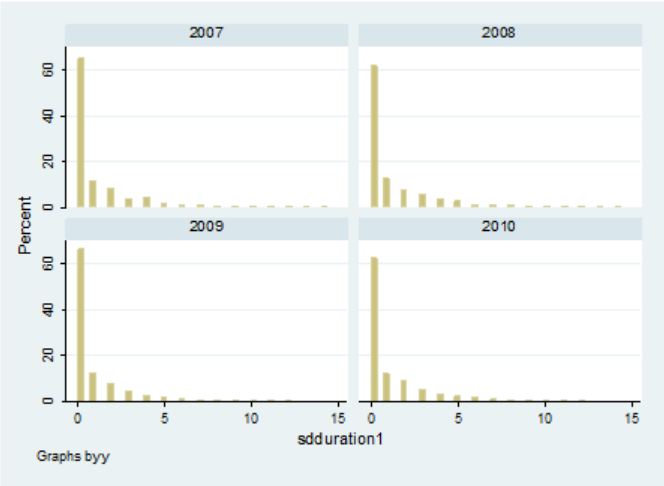

Figure 3: Distribution of the Timing to Reconnect

What we find is that the majority of self-disconnections (>62%) last for less than one day (with households topping-up their meter on the same day they are predicted to have run out of credit); between 72% and 82% of self-disconnections last for less than two days. 12% to 18% of self-disconnections last longer than 3 days.

3.6 The Extent of Self-Disconnection Revisited 3

Finally, we can use our metering data to relate the frequency with which households self-disconnect and the duration they are without electricity. Figure 4 below shows this relationship. (It is based on a local linear regression – which allows us to smooth the data).

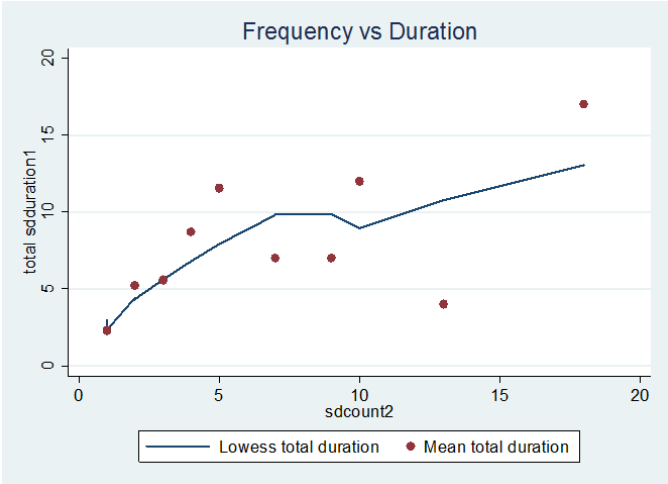

Figure 4: Linking Frequency and Duration w/o Elec.

What we find is a positive relationship between the two variables – such that households which self-disconnect once per year tend to be without electricity for approximately one day per year while households which self-disconnect five times per year tend to be without electricity for almost seven days per year. This suggests that – at least for some households – self-disconnection constitutes a significant burden.

4 Self-Disconnection: A Hard Policy Problem

The previous literature (see Table 2) has identified four possible drivers of self-disconnection:

1. Financial Constraints;
2. Forgetting to top-up one's meter;
3. Coordination issues among household members when it comes to purchasing credit; and
4. Unavailability of credit (e.g. because of a closed outlet)

In this section, we test the relative importance of the four potential drivers using our metering data.

4.1 The drivers of self-disconnection revisited

One way to test why households self-disconnect is by regressing how often households have self-disconnected in 2010 on measures/proxies for each of the four potential drivers.

That is, we can regress the number of self-disconnections on a measure of financial constraints; a self-stated measure of forgetfulness; a measure of who is responsible for purchasing credit[13] (as a proxy for coordination issues); a measure of how much opening hours of outlets selling top-up are perceived as an issue (as a proxy for unavailability of credit); and a set of controls .

Our model then takes the following form:

$$SD_i = \gamma_o + \gamma_1 X_i + \gamma_2 Z_i + \varepsilon_i \ (1)$$

where SD_i is the number of self-disconnections of household i in 2010. X_i is a vector summarising our main variables (measuring financial constraints, forgetfulness etc); Z_i a vector of control variables and γ the corresponding regression coefficients. ε_i is an error term. We estimate (1) by means of OLS (using a log-lin specification)[14] and by means of Generalised Methods of Moments.[15]

4.2 A Note on the Data

For the analysis, we match our metering data with data from a telephone survey (same as before).[16] The households in our survey were randomly drawn from the population of all households with an electricity pre-payment account with British Gas. Within each household, our interview was directed at the household member 'in charge of purchasing top-up'.

To deal with the large number of households that have not self-disconnected/ have self-disconnected only once in 2010 we stratified our sample along the number of self-disconnections – giving more weight to households which have

[13] The answer possibilities ranged from: always the same person to never the same person.

[14] That is, we use the log of the number of self-disconnections in 2010 as our dependent variable; code all zero counts as 'ones' – and regress the log of self-disconnections on our key variables and a dummy variable that equals one for all zero counts. See: Hausman, Hall and Griliches (1984).

[15] Both approaches allow us to instrument for possible problems of measurement error; zero inflation will be addressed in the next the section. In case of our GMM estimation, we specified both an additive and multiplicative error term (as proposed by Mullahy, 1997). The results, however, were virtually identical.

[16] The fieldwork was carried out by Ipsos Mori.

self-disconnected more often. Table 4 provides summary statistics of the stratification we used.[17]

	Strata 1	Strata 2	Strata 3	Strata 4
Number of Self-Disconnections	0	1-2	3-5	6+
Number of observations	126	125	126	125
Share in population	78%	16.5%	3.85%	1.65%
Share in sample	25%	25%	25%	25%

Table 4: Overview Stratification

To capture whether households are 'financially constrained' we asked them whether they have been in arrears (with regard to mortgage, rent, loan instalments etc) in the last year; whether they have been assisted with debt repayment by a debt advice agency; and how they regard their payment capacity to cover bills (very good to very difficult)?

Following the European Commission (2008), we code a household as 'financially constrained' if it reports difficulties with respect to either of these dimensions. To capture 'forgetfulness' of the household member 'in charge of purchasing top-up', we asked him/her in the interview whether he/she considers him-/herself forgetful.[18] Summary statistics of our key variables are provided in Appendix A.

4.3 Analysis

Table 5 provides the results from our estimation of (1).[19]

The first column and third column show our baseline results. The second column and fourth column show the results from the same analysis if we instrument for 'forgetfulness' – to account for possible problems of measurement error.

[17] The response rate was: 60%.

[18] See e.g. Commissaris and Ponds (1998)

[19] Unweighted results. To guard against the possibility that our inference is unduly affected by sampling weights, we estimated our model also with sampling weights. What we find is that our point-estimates tend to be smaller; the same is true for the standard errors with sampling weights - but the main conclusions remain the same.

	OLS Estimation		GMM Estimation	
	(1)	(2)	(3)	(4)
Financial Constraints	0.19***	0.20***	0.35***	0.55*
	(0.06)	(0.08)	0.05	(0.35)
Forgetfulness	0.07	-0.11	0.08	-0.13
	(0.08)	(1.05)	0.06	(1.05)
Reachabililty	0.02	0.03	0.004	-0.14
	(0.10)	(0.11)	0.08	(0.39)
Coordination Issues	0.07	0.08	0.13**	0.01
	(0.07)	(0.08)	0.06	(0.34)
Socio-economic Controls	Yes[20]	Yes	Yes	Yes

Table 5: Results Regression Analysis of (1) - Statistically Significant at 1% ***; 5% **; 10% *

What we find is that only one variable is statistically significantly related to the number of self-disconnections: which is, 'financial constraints'. Specifically, we find that being 'financially constrained' is associated with an increase in the average number of self-disconnections by between 19 and 35 percent.

This result is robust (qualitatively and quantitatively) to the estimation method we use. It is also robust to instrumenting for 'forgetfulness' (using 'age of the household member in charge' as an instrument).

4.4 Robustness Check

One possible explanation why we find a statistically significant effect for 'financial constraints' but none of the other potential drivers of self-disconnection is that the other drivers have a significant effect on the number of self-disconnections only at certain points in the distribution of self-disconnections.

To give an example, it is possible that households in which the household member 'in charge of purchasing top-up' is 'forgetful' are either effective in dealing with this problem or not effective at all (but nothing in between).

In this case, we would expect that 'forgetful households' are no more likely to self-disconnect once than other households, but – given that they have self-disconnected once – are significanlty more likely to self-disconnect again (and again).[21]

[20] Unless otherwise stated, in the following 'socio-economic controls' include: age; gender; marital status; labour market status and housing. To estimate (1), we also included a control variable capturing whether a household has a car/van - to control for differences in mobility.

[21] See e.g. Agarwal et al (2008)

What this means is that the average effect of forgetfulness may be insignificant even though forgetfulness has a significant effect on the probability that households self-disconnect at a certain point in the distribution of self-disconnections. One way to check if this is the case is by using a quantile regression approach.

Unlike OLS (or GMM)[22], a quantile regression is not limited to explaining the mean of the number of self-disconnections, but allows us to explain the determinants at any point of the distribution of self-disconnections.[23] The results from a quantile analysis of (1) are presented in Figures 5a-5d below.[24]

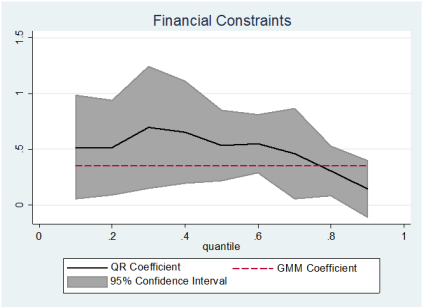

Figure 5a: QR Financial Constraints

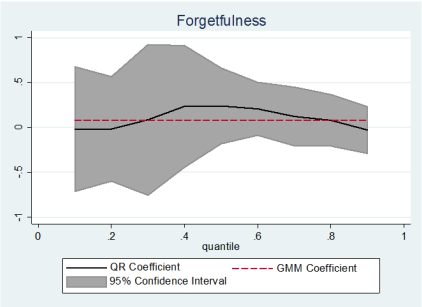

Figure 5b: QR Forgetfulness

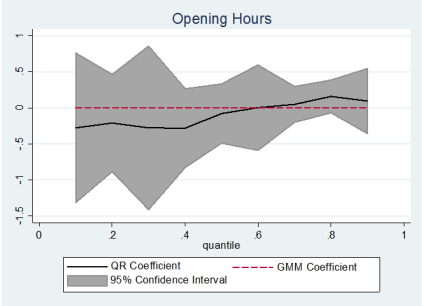

Figure 5c: QR Opening Hours

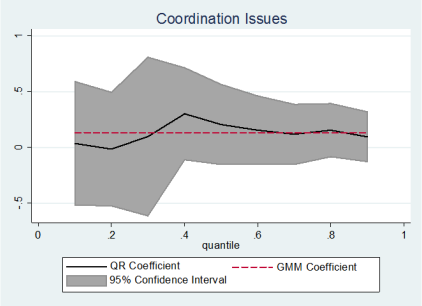

Figure 5d: QR Coordination Issues

The horizontal lines represent the GMM point estimates of each coefficient from before. The moving lines plot how these estimates vary over different quantiles of our dependent variable (and the corresponding confidence intervals).[25]

[22] By that we mean standard GMM. There are quantile GMM approaches which are not limited to explaining the mean of a dependent variable. See e.g. Abadie (2003); or Lee (2007)..

[23] See Koenker, Bassett (1978); and Koenker, Hallock (2001)

[24] We estimate (1) using Machado and Silva's (2005) 'jittering approach'.

[25] A simple Hausman-Wu test suggests that we cannot reject the null-hypothesis of no difference between our point estimate for forgetfulness (in columns 3 and 4 in Table 5). This is why we do not instrument for 'forgetfulness' in this section.

In the top left figure, for example, it shows that the effect of being 'financially constrained' has a larger effect on the number of self-disconnections (by ca 150%) in the top quantiles than the bottom quantiles. The same is not true if we look at the top right figure of Figure 5. It shows the impact estimate of 'forgetfulness' – which shows a peak in the middle quantiles.

The bottom left and right figures show that 'coordination issues' among household members and 'unavailability of credit' remain insignificant (and close to zero) across all quantiles of self-disconnections. This suggests that 'coordination issues' and 'unavailability of credit' are indeed of secondary importance for why households self-disconnect.

5 Self-Disconnection: An Alternative Perspective

In the last section, we showed that the main driver of self-disconnection is 'financial constraints' (and to a lesser extent 'forgetfulness').

In this section, we argue that – even if this suggests that it is likely to be hard/expensive to reduce the total number of self-disconnections – it may still be possible to reduce the negative impact of self-disconnections (at least to some extent) by helping households to better smooth the number of self-disconnections over the course of a year.

5.1 When should a household self-disconnect (if it has to do so)?

The starting point for our discussion is the idea that – from a rational perspective – households should self-disconnect whenever it hurts them the least.

What this suggests is that (in the absence of seasonal differences in prices and/or tastes) households should spread the total amount/total duration of self-disconnection evenly over the course of a year – that is, even if they face a strong seasonal pattern in income flows/energy use. In the latter case they would simply borrow/save to ensure that at any point they have the right amount of cash available.

To see this more formally, suppose that all a household cares about is its energy use (electricity + gas). Further, assume that a household requires a fixed amount of credit to avoid self-disconnection every month/season; that this

amount is the same every month/season – and that the only thing that varies across months/seasons is disposable income.

One way to think about these assumptions is as follows: once we take into account the higher costs for energy in the autumn/winter, households – with otherwise constant incomes over the course of a year[26] – are exposed to varying disposable incomes to purchase a fixed amount of top-up (with higher incomes in the spring/summer and lower incomes in the autumn/winter).

Given these assumptions: how much credit will a household demand each month/season? Or: how much will it self-disconnect? We can analyse these problems – without loss of generality – using a simple two-season framework (see e.g. Paxson, 1993) in which households maximise additively separable utility such that:

$$\max \sum_{t=0}^{\infty} \rho^{2t}[U(C_{oit};\alpha_0) + \rho U(C_{1it};\alpha_1)] \quad (3)$$

$$\text{subject to } \sum_{t=0}^{\infty} R^{-2t}(P_0 C_{0it} + P_1 \tfrac{C_{1it}}{R}) = W_i + \sum_{t=0}^{\infty} R^{-2t}(Y_{0i} + \tfrac{Y_{1i}}{R}) \quad (4)$$

where C_{jit} is demand for credit of household i in season j in year t, P_j represents the price of credit in season j in all years, W_i is initial financial wealth, and ρ is a seasonal discount rate. The term α_j is a season specific taste parameter, which is assumed (for now) to be identical across households, and Y_{ji} is disposable income of household i in season j.[27]

If it is assumed that $\rho R = 1$, then utility maximisation yields two season-specific levels of demand for credit for each household – C^*_{0i} and C^*_{1i} – that do not vary across years. To derive closed-form expressions for consumption in each season, we assume that utility has a constant relative risk aversion (CRRA) form with a risk aversion parameter of a, such that $U`(C_{ji};\alpha_j) = \alpha_j (C_{ji})^{-a}$.

This assumption, together with the assumptions noted above, yields the following expressions for demand for credit in each season:

$$C^*_{0i} = \frac{\lambda R}{P_o \lambda R + P_1}[Y_{0i} + \frac{Y_{1i}}{R} + W_i(\frac{R^2-1}{R^2})] \quad (5)$$

$$C^*_{1i} = \frac{R}{P_o \lambda R + P_1}[Y_{0i} + \frac{Y_{1i}}{R} + W_i(\frac{R^2-1}{R^2})] \quad (6)$$

where $\lambda = (\frac{\alpha_1 P_0}{\alpha_0 P_1})^{\frac{-1}{a}}$

[26] A simple analysis of the corresponding data from the LCF suggests that this assumption is fair – irrespective of whether we look at 'disposable income' or 'total household expenditure'.

[27] We also assume that individuals can borrow and save at a constant seasonal interest rate r=R-1. We will come back to this assumption later.

In each of the equations, the term within brackets is the same and represents 'permanent income'. Given that in our model self-disconnection is merely the flip side of consumption – it is the (negative) difference between the demand for credit and the fixed need to avoid self-disconnection – and (for now) that the duration of self-disconnections is constant over time, the two expressions confirm our basic intuition:

1. In the absence of differences in prices and/or tastes across seasons rational households spread the number of self-disconnections evenly over the course of a year and

2. Although income and wealth affect the number of self-disconnections in each season, they do not affect how these self-disconnections are allocated across seasons. That is, the timing of self-disconnections is independent of the timing of income flows/energy use.[28]

5.2 Sub-optimal Self-Disconnection Behaviour?

A preliminary test of whether households do behave optimally – and smooth the number of self-disconnections evenly over the course of a year – is by plotting the probability that households self-disconnect for each month of the year (and check if it is constant over time).

Figure 6a shows the result from this analysis[29] (in which we estimated the probabilities to self-disconnect by regressing whether a household self-disconnects in a particular month on a set of month dummies, year dummies and household fixed effects).

[28] One short-coming of our (one good) model is that it allows for households to demand/use more credit in a given season than is necessary to avoid self-disconnection (which is implausible). To the extent, however, that what we are interested in is situations in which the optimal demand for credit falls below the minimum necessary to avoid self-disconnection this property has no bearing on the main insights of our model. We will discuss corner solutions later on.

[29] We use only households which have self-disconnected at least once in the past 4 years. (So all estimated probabilities (relative to December) should be considered conditional in this sense).

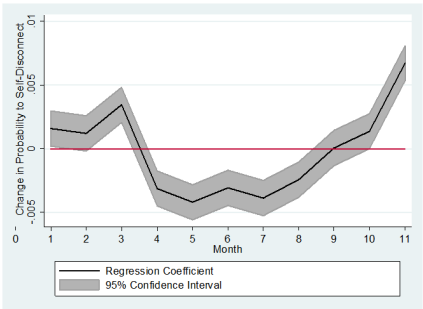

Figure 6a: Regression Coefficients – Self-disconnection

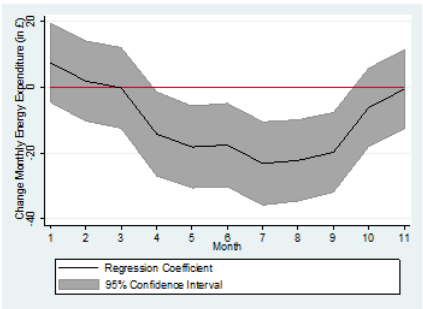

Figure 6b: Regression Coefficients – Electricity + Gas expenditure

What we find is that – contrary to the prediction of our theoretical model – there is a strong seasonal pattern in the probability that households self-disconnect in a particular month (relative to the probability that they self-disconnect in December): A simple F-test suggests that we can reject the null-hypothesis of 'no seasonal pattern' at the 1% level of significance.[30]

A preliminary test of whether the second implication of our model stands up to the data is to compare the seasonal pattern in self-disconnection with the seasonal pattern in energy expenditure. Figure 6b does exactly this: it plots the coefficients from a regression of households' total gas and electricity expenditure in each month on a set of month dummies, year dummies and a set of controls.[31]

What we find is a similar seasonal pattern in gas and electricity expenditure as for the probability that households self-disconnect. In particular, we find the same local minima and maxima – and spikes in March; April and July (albeit less pronounced in the energy data).[32] This suggests that – again contrary to the prediction of our simple model – the timing of self-disconnections may not be entirely independent of the timing of households' income flows/energy use.

5.3 Is there are real problem?

The conflict between the optimal self-disconnection behaviour (as suggested by our simple model) and the way people actually behave (as reflected in the data) rests on two critical assumptions:

[30] The F-statistic is 43.98. The corresponding p-value: <0.01

[31] For this analysis, we use data from the LCF/Expenditure and Food Survey (EFS) from 2007-2009. It was not possible to extract and match the corresponding metering information.

[32] The most likely reason for the less pronounced seasonal pattern is measurement error in the survey data (and the relatively small sample size).

- The first one is that there are no seasonal differences in the price for credit and/or tastes.[33]
- The second assumption is that self-disconnections last for the same period of time over the course of a year.

If the price of credit is higher in the autumn/winter and/or households prefer to self-disconnect in the autumn/winter to self-disconnecting in the spring/ summer, it can be optimal for households to not smooth their self-disconnections evenly over the course of a year (but delay some of their self-disconnections until later in the year).[34]

Similarly in case of the second assumption: if self-disconnections last longer in the spring/summer than the autumn/winter, it is possible that we find a strong seasonal pattern in the probability that households self-disconnect – even if households do in fact optimally smooth the total amount/ total duration of self-disconnection over the course of a year.

What this suggests is that our simple plots in the last section are not sufficient to establish that households top-up behaviour is sub-optimal. (After all, even the co-movement with households' income flows/energy use may simply be accidental).

In addition, it suggests that to actually establish that the pattern in self-disconnection we find is sub-optimal what we need to do is: show that it is not driven by seasonal differences in prices; tastes and/or the duration of self-disconnections (but by the seasonal pattern in income flows/energy use instead).

5.4 Empirical Framework

One way of testing whether the seasonal pattern in self-disconnections in the last section is due to seasonal differences in prices; tastes and/or the duration of self-disconnections or a seasonal pattern in income/energy use is by comparing the probabilities that households self-disconnect in the autumn/winter when households face a strong seasonal variation in income/energy use and when they do not.

[33] Another possibility – which we do not consider in our model – is uncertainty. The evidence in the next section, however, is at odds with this idea.

[34] Formally, the idea is that if $P_1 > P_0$ and/or $\alpha_0 > \alpha_1$, then $\lambda = (\alpha_1 P_0/\alpha_0 P_1)^{-1/a} > 1$ which in turn implies that $C^*_{0i} = \frac{\lambda R}{P_o \lambda R + P_1}[Y_{0i} + \frac{Y_{1i}}{R} + W_i(\frac{R^2-1}{R^2})] > C^*_{1i} = \frac{R}{P_o \lambda R + P_1}[Y_{0i} + \frac{Y_{1i}}{R} + W_i(\frac{R^2-1}{R^2})]$ (i.e. that consumption in the spring/summer is smaller than consumption in the autumn/winter) which, in turn, suggests that the number of self-disconnections is higher in the autumn/winter than in the spring/summer.

The idea is that if the seasonal pattern in self-disconnections is driven by variation in prices; tastes and/or the duration of self-disconnections, we would expect to find the same probabilities under both scenarios. The same is not true if the seasonal pattern in self-disconnections is due to seasonal differences in income/energy use. In this case, we would expect a higher probability to self-disconnect under the first scenario than the second one.

Unfortunately, we do not observe the same households exposed to seasonal changes in income/energy use and at the same time not exposed to seasonal changes in income/energy use. One way of dealing with this is to compare the probabilities to self-disconnect in the autumn/winter of households which are exposed to seasonal changes in income/energy use and households which are not (or much less) exposed to such changes.

The problem with this is that the difference in probabilities between the two groups of households may be due to reasons other than differences in exposure to seasonal variation in income/energy use – such as differences in household income, family size, geographical location etc.[35] One way to deal with this is by comparing the changes in the probabilities to self-disconnect from one season to the next (rather than levels).

If the reason for the pattern in self-disconnections is due to differences in prices; tastes and/or the duration of self-disconnections – there is little reason to expect that the probability to self-disconnect increases more in the autumn/winter for households which are more exposed to seasonal changes in income/energy use than for households which are less exposed to such changes – even if the two groups differ in terms of their background characteristics).

If, on the other hand, the seasonal pattern in self-disconnections is due to seasonal differences in income/energy use, we would expect to find a stronger increase in the probability to self-disconnect in the autumn/winter for households which are more exposed to seasonal changes in income/energy use than for households which are less exposed to such changes.

5.5 Analysis

One way to implement the test outlined in the last section is by comparing the changes in the probabilies to self-disconnect (with an electricity pre-payment meter) of households which have a gas pre-payment meter and households paying

[35] As described in section 2.2.

for their gas by means of a budgeting scheme.[36]

The idea is that: while expenditure on gas is highly seasonal for households using a gas pre-payment meter, the same is not true for households paying for their gas by means of a budgeting scheme. Figure 7a and 7b below illustrate this:

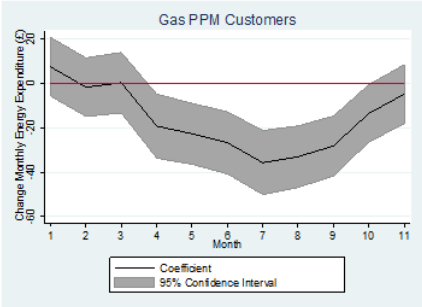

Figure 7a: Energy Expenditure Gas-PPM

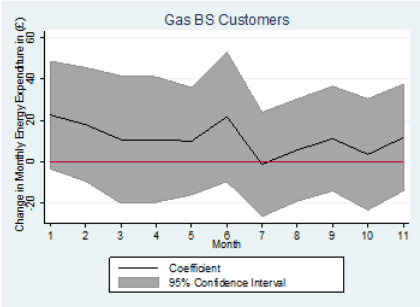

Figure 7b: Energy Expenditure Gas-Budgeting Scheme

The figures show the monthly changes in total gas and electricity expenditure over the course of a year (relative to December) – separately for household with a gas pre-payment meter and households on a budgeting scheme.[37] They show that – in line with what we would expect – energy use is highly seasonal for customers with a gas pre-payment meter, while it is relatively flat for households paying for their gas by means of a budgeting scheme.

We can now test whether the seasonal pattern in self-disconnection is different between the two groups of households: Figures 8a and 8b do exactly this: they plot the changes in the probability that households self-disconnect (with their electricity pre-payment meter) for each month of the year – separately for households with a gas pre-payment meter and households paying for their gas by means of a budgeting scheme.[38]

[36] Budgeting scheme includes: quarterly equal payment; monthly direct debit (same amount each month); and mag card customers (with a set amount each week /fortnight/ month).

[37] The corresponding regression equation and quantitative estimates are provided in Appendix B.

[38] The corresponding regression equation and quantitative estimates are provided in Appendix B - again.

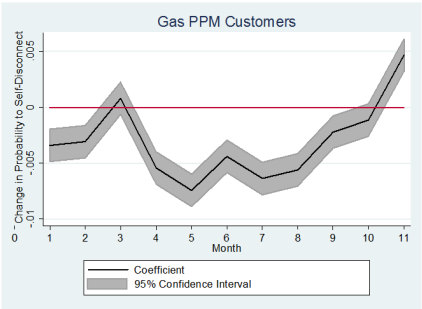

Figure 8a: Self-Disconnection - Gas-PPM

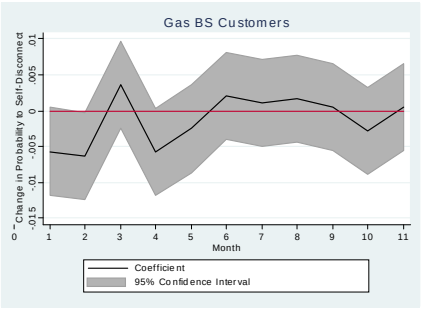

Figure 8b: Self–Disconnection - Gas Budgeting Scheme

What we find is that – in line with the idea that the seasonal pattern in self-disconnection is driven (at least to some extent) by the seasonal pattern in income/energy use – households with a gas pre-payment meter exhibit a much stronger seasonal pattern in self-disconnection with their electricity pre-payment meter than households paying for their gas by means of a budgeting scheme. A simple F-test shows that we can reject the null-hypothesis of identical seasonal patterns at the 1% level of significance.[39]

The identification assumption of our test is (of course) that in the absence of the stronger seasonal pattern in income/energy use, the changes in the probability to self-disconnect from spring/summer to autumn/winter would have been the same for the two groups. In Appendix C, we provide a series of tests to explore how plausible this assumption is. We come to the conclusion that it is plausible and hence that our main test is meaningful.

6 Why do households fail to smooth SDs?

In the last section, we showed that households exhibit a strong seasonal pattern in self-disconnections. In addition, we showed that this pattern is sub-optimal by showing that it is not driven by seasonal differences in prices; tastes and/or the duration of self-disconnections (but seasonal differences in income flows/energy use instead).

This suggests that we might be able to improve the self-disconnection situation of households – by helping them to better smooth the number of self-disconnections over the course of a year. In this section, we make a first

[39] Detailed quantitative estimates can again be found in Appendix B.

step in this direction by exploring why households fail to smooth their self-disconnections.

6.1 A possible explanation: Preference Reversals

One possible explanation why many households fail to smooth the number of self-disconnections over the course of a year is: preference reversals.

The way preference reversals can affect households' ability to smooth self-disconnections is by affecting their ability to save: if households fail to save in the spring/summer, they will be more exposed to the seasonal dip in disposable income/seasonal increase in energy use in the autumn/winter – which, in turn, affects their probability to self-disconnect.[40]

The basic idea of preference reversals is that from a long-run perspective, individuals with preference reversals tend to have one set of preferences but when the future arrives, tend to have a different set of preferences. It is straightforward to see how this 'time-inconsistency' in preferences can affect households' ability to save in the spring/summer:

While in the autumn/winter households might have a strong preference to save in the spring/summer to avoid (excessive) self-disconnection in the following autumn/winter, as soon as the spring/summer arrives their preferences change and – instead of saving – they keep on consuming (making themselves vulnerable to self-disconnection in the following autumn/winter).

There are several possible reasons why individuals may exhibit preference reversals. One is changes in the saliency of costs and benefits (Akerlof, 1991): the idea is that in the autumn/winter the benefits from saving during the spring/summer loom large, while the costs seem small. But as the spring/summer approaches, individuals become increasingly aware of the costs of saving, while the benefits become increasingly less clear.

Preference reversals have been shown to be at the core of a series of economic puzzles – ranging from why students tend to start too late to prepare for their exams (Ariely and Wertenbroch, 2001); to why individuals often save too little for their retirement (Laibson, 1997); to why they tend to fail to go to the gym as often as they would like to (Della Vigna and Malmendier, 2005); to why they tend to eat more than is good for them (Scharff, 2009).[41]

[40] This assumes that households which self-disconnnect (frequently) tend to be borrowing constrained - which seems reasonable given the high correlation between 'financial constraints' and the number of self-disconnections (see section 4.3).

[41] Basu and Wong (2009) provide a theoretical model which incorporates preference reversals

6.2 Testing for Preference Reversals

The simplest way to test whether preference reversals (also) affect households' ability to smooth self-disconnections over the course of a year is by regressing the difference in the number of self-disconnections in the autumn/winter and spring/summer on a measure of preference reversals.

If the seasonal pattern in the probability that households self-disconnect is driven by preference reversals, we would expect that households in which the household member in charge of purchasing top-up suffers from preference reversals are more likely to self-disconnect primarily in the autumn/winter.

The corresponding model takes the following form:

$$DSD_i = \gamma_0 + \gamma_1 PR_i + \gamma_2 Z_i + \zeta_i \ (7)$$

where DSD_i is the difference (in logs) of self-disconnections in the autumn/winter and spring/summer of household i in the years 2007-2010.[42] PR_i is a variable capturing preference reversals (of the household member who is in charge of purchasing top-up); Z_i a vector of control variables and γ the corresponding regression coefficients. ζ_i is an error term.

6.3 A Note on the data

As is common in the literature, we measure preference reversals by:

1. Asking individuals (in charge of purchasing top-up) to choose between receiving a smaller reward immediately and receiving a larger reward with some delay and then

2. Asking the same question with the same rewards but at a further time frame (see Tversky and Kahneman, 1986; Benzion et al, 1989).

Our sample questions are as follows: would you prefer to receive £350 guaranteed today or £400 guaranteed in 1 month? and: would you prefer to receive £350 guaranteed in 6 months or £400 guaranteed in 7 months?[43]

in a seasonal model of consumption (similar to the one we present in section 5.1). They show that – analytically – this simple extension entails a strong correlation between the timing of income flows and the timing of consumption (i.e. in our case, between the timing of income flows/energy use and the timing of demand for credit/the timing of self-disconnections).

[42] We calculate an average number of self-disconnections for the spring/summer months and autumn/winter months before taking the difference (where the averages take into account the number of spring/summer months and autumn/winter months a household had its prepayment meter).

[43] In both cases, if respondents chose £350, we asked them the same question with £450 (rather than £400) as alternative amount. If respondents still chose £350, we asked them

Following Asharf et al (2006), we call the first question the 'near-term' frame and the second question the 'distant' frame choice. We interpret the choice of the immediate reward in either of the frames as 'impatient'.

If an individual chooses the immediate reward in the near-term frame and the delayed reward in the distant frame we say he/she displays preference reversals. In particular, we refer to this case as 'hyperbolic' (since the implied discount rate in the near term frame is higher than that of the distant frame).

Table 6 below describes the cell densities for the different possible outcomes.

			Indifferent between 350 in 6m & X in 7m			
			Patient	Sw Impat.	M impat.	Total
			X<350	350<X<400	X>450	
Indiff.	Patient	X<400	250	45	29	324
bw 350	Sw Impat.	400<X<450	35	21	12	68
now &	M impat.	X>450	57	11	42	110
X in 1m	Total		342	77	83	502

Table 6: Summary statistics Hypothetical Choice Experiment

The table shows that approximately 21 percent of individuals in our sample are hyperbolic, that is, these individuals are more patient over future trade-offs than current trade-offs; 16 percent show preference reversals in the opposite direction (being less patient over future trade-offs than current ones); and 63 percent of individuals show no preference reversals.

6.4 Analysis

Table 7 shows the results from our estimation of (7) – linking the timing of self-disconnection and our measure of (hyperbolic) preference reversals.[44]

'how much would we have to offer you to take the higher amount'. Please note: there was a break of approx. 15 minutes between the 'near-term' frame and the 'distant' frame.

[44] We keep only households in sample which have self-disconnected at least twice in the last 4 years – since only this way we can infer something about their ability to smooth consumption/self-disconnections from their self-disconnection behaviour. We delete 122 observations in the process. As before, we estimate our model with and without survey weights. We find that, again, our main conclusions do not change if we use survey weights. In fact, our point estimates are almost identical.

	OLS w/o controls	OLS with controls
Hyperbolic	0.41	0.39**
	(0.18)	(0.18)
Impatient now vs 1 m	-0.39	-0.41**
	(0.17)	(0.18)
Distance from Nearest Outlet	-0.03	-0.03
	(0.07)	(0.07)
Socio-economic Controls	No	Yes

Table 7: Regression Analysis of (2) - Stat

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

What we find is that – in line with our hypothesis – households in which the individual in charge of purchasing top-up suffers from (hyperbolic) preferences reversals are more likely to self-disconnect primarily in the autumn/winter than households in which the person in charge does not suffer from (hyperbolic) preference reversals.

Specifically, what we find is that suffering from (hyperbolic) preference reversals (by the household member in charge of purchasing top-up) is associated with a statistically significant 40% increase in the ratio of the number of self-disconnections in the autumn/winter compared to the spring/summer.

This finding is robust to different specifications of (7). In addition, we can rule out several alternative explanations for our finding – such as noise; an inability to understand our survey questions; the timing of income flows – as documented in Appendix D.

7 Exploring Different Policy Responses

In the last section, we provided evidence which suggests that one important driver of households' sub-optimal self-disconnection behaviour is preference reversals.

In this section, we discuss different ways to address this problem. We focus on two broad classes of policies which can help to mediate the negative effect of preference reversals on households' ability to smooth their self-disconnections. They are:

- An increase in awareness of preference reversals in the context of self-disconnections and

- Providing households (which are aware of the problem) with a commitment savings device to 'lock in' their optimal behaviour.

The idea of the first class of policies is that if household members are aware of their preference reversals, they can try to commit themselves irreversibly in the autumn/winter to a particular behaviour in the spring/summer to so avoid deviating from their (long-term) plans.

This can be done either by the household member in charge of purchasing top-up – in which case we speak of 'sophistication' as a mediating factor (see e.g. O'Donoghue and Rabin, 2001) or by his/her partner – in which case we speak of 'interference' as a mediating factor (see e.g. Ashraf, 2009 or Brutscher, 2011).

In both cases, however, we need that household members have the possibility to 'lock in' their behaviour and so prevent either themselves or their partners from spending their savings before the autumn/winter. One way households can do this is by means of a comitment savings device:

A commitment savings device is a savings device which restricts access to one's savings – to certain times of the year (e.g. the autumn/winter) and/or to certain uses of one's savings (e.g. the purchase of top-up) and so allows one to 'lock in' one's behaviour.[45]

7.1 A Two Step Approach

One way of exploring whether the main challenge for policy makers is to increase awareness about preference reversals in the context of self-disconnections or to (also) provide households with a commitment savings device is by testing whether 'sophistication' and/or 'interference' has/have a mediating effect on preference reversals.

The idea is that if we find that 'sophistication' and/or 'interference' has/have a mediating effect, this tells us that the main challenge for policy makers is to raise awareness among those households which are not aware of the negative role of preference reversals yet (when it comes to smoothing self-disconnections over the course of a year).

[45] An example of a commitment savings device – in a context similar to ours – is fuel stamps in Northern Ireland. Fuel stamps are small pieces of paper which can be purchased at specified outlets (in Northern Ireland), collected on fuel stamps savings cards, and used to pay for one's oil bill in the autumn/winter. Because the stamps are non-refundable and can only be used for the purchase of heating oil, they allow household members to prevent themselves/their partners from spending their savings on anything else but heating oil. For more details on fuel stamps see Brutscher (2011)

If, on the other hand, we find that even households which are 'sophisticated' and/or satisfy the conditions for 'interference' fail to deal with their preference reversals in an effective way, this suggests that the main challenge for policy makers is to provide households with a commitment savings device – possibly on top of increasing awareness among households which are (still) 'unaware'.

One way to implement this test is by dividing households into several different groups – depending on their level of 'sophistication' and 'interference' – and to then estimate (7) separately for these groups of households. The idea is to so check whether the effect of preference reversals is (significantly) stronger/weaker for any of the groups of households. There are, however, several problems with this approach:

First, there may be more mediating factors in our data than 'sophistication' and/or 'interference' – like gender; age; etc – which we would like to capture. Second, even if 'sophistication' and/or 'interference' were the only two such factors, the fact that we have more than one (potential) mediating factor makes it difficult to separate the data into groups of households without losing a lot of explanatory power.[46]

Finally, although the measures of 'sophistication' and 'interference' we use for our analysis are the best ones available to us, they are likely to not fully capture the two concepts. This makes it (again) difficult to separate our data into groups of households without losing a lot of explanatory power in the process.

A more promising approach to test for the mediating effect of 'sophistication' and/or 'interference' employs a two step approach:

1. In the first step, we use a finite mixture approach to divide the sample into two groups of households and to estimate (7) separately for both groups

2. In the second step, we then test whether the two groups differ in terms of 'sophistication' and/ or 'interference'.

The idea of the first step is to (simultaneously) divide households into two groups of households and test whether preference reversals have a different effect for group 1 and group 2. The idea of the second step is – given that preference reversals have a different effect in the two groups – to check what mediates the effect in one of the two groups (but not the other).

This two-step approach has the advantage that it can account for more than the two (potentially) mediating factors which we have identified; it allows us to

[46] What makes this problem worse is that one of our mediating factors is multi-categorical. This means that we would have to divide our sample into even more sub-samples.

test for 'sophistication' and 'interference' using the largest possible sample size; and it minimises the problem of noise in our measures of 'sophistication' and/or 'interference'.[47]

7.2 First Step

The main challenge of our approach is how to divide households into two groups (in the first step). We deal with this by means of a 'statistical trick':

The idea is that we can get an increasingly precise idea of group membership by going through an iterative procedure in which we choose the parameters of our model and group membership so as to maximise the fit of our model.[48]

[47] Regarding the last point, please note that our point estimates of the mediating effect of sophistication and/or interference does not depend on our measures of these variables.

[48] Consider Do and Batzoglou's (2008) simple coin-flipping experiment in which we are given a pair of coins A and B of unknown biases, Xa and Xb, respectively (that is, on any given flip, coin A will land on heads with probability Xa and tails with probability 1-Xa and similarly for coin B). Our goal is to estimate X=(Xa, Xb) by repeating the following procedure five times: randomly choose one of the two coins (with equal probability), and perform ten independent coin tosses with the selected coin. Thus the entire procedure involves a total of 50 coin tosses.

During our experiment, suppose that we keep track of two vectors x=(x1,x2...) and z(z1, z2...) where x is the number of heads observed during the ith set of tosses, and z is the identity of the coin used during the ith set of tosses. Parameter estimation in this setting is known as the complete data case in that the values of all relevant random variables in our model (that is, the result of each coin flip and the type of coin used for each flip) are known. We estimate Xa and Xb by means of ML (roughly speaking, the ML method assess the quality of a statistical model based on the probability it assigns to the observed data).

Now consider a more challenging variant of the parameter estimation problem in which we are given the recorded head counts x but not the identifies z of the coins used for each set of tosses. We refer to z as hidden variables or latent factors. Parameter estimation in this new setting is known as the incomplete data case. This time, computing proportions of heads for each coin is no longer possible, because we don't know the coin used for each set of tosses. However, if we had some way of completing the data (in our case guessing correctly which coin was used in each of the five sets), the we could reduce parameter estimation for this problem with incomplete data to maximum likelihood estimation with complete data.

One iterative scheme for obtaining completions could work as follows: starting from some initial parameters X=(Xa, Xb) determine for each of the five sets whether coin A or coin B was more likely to have generated the observed flips (using the current parameter estimates). Then, assume these completions (that is guessed coin assignments) to be correct, and apply the regular maximum likelihood estimation procure to get Xt+1. Finally, repeat these two steps until convergence. As the estimated model improves so too will the quality of the resulting completions.

The expectation maximisation algorithm is a refinement on this basic idea. Rather than picking the single most likely completion of the missing coin assignments on each iteration, the expectation maximisation algorithm computes probabilities for each possible completion of the missing data, using the current parameters X. These probabilities are used to create a weighted training set consisting of all possible completion of the data. Finally, a modified version of ML estimation that deals with weighted training examples provides new parameter estimates Xt+1. By using weighted training examples rather than choosing the single best completion, the expectation maximisation algorithm accounts for the confidence of the model in each completion of the data. In summary, the EM algorithm alternates between the steps of guessing a probability distribution over completions of missing data given the current model

The way we implement this approach formally is as a mixture of two normal regressions or switching regressions:

$$Y_i = X_i\beta_1 + \varepsilon_{1i} \text{ with probability } \theta \ (8)$$

$$Y_i = X_i\beta_2 + \varepsilon_{2i} \text{ with probability } (1-\theta) \ (9)$$

where ε_{1i} and ε_{2i} are mutually independent, *iid* normal with zero means, and variance σ_1^2 and σ_2^2, respectively.

7.3 Analysis

The results of estimating (8) and (9) – using an Expectation Maximisation Algorithm[49] – are contained in columns 2 and 3 of Table 8.

	OLS Spec.	FM Group1	FM Group2
Sample Share	X	0.69	0.31
Hyperbolic	0.39**	0.25	0.89***
	(0.39)	(0.22)	(0.28)
Impatient now vs 1 m	-0.41	0.16	-1.75***
	(0.18)	(0.21)	(0.60)
Distance from Nearest Outlet	-0.03	-0.10	-0.09
	(0.07)	(0.09)	(0.14)
Socio-economic Controls	Yes	Yes	Yes
Σ	0.73	0.50	0.46

Table 8: Mixture Analysis - Statistically Significant
at 1% ***; 5% **; 10% *

The table shows that about two thirds of the households in our sample are associated with group 1 (69%) and about one-third (31%) with group 2. In addition, the table shows that the estimated standard error[50] of the two groups are both (0.50 and 0.46) smaller than their OLS counterpart (0.73).

This suggests that – in line with the idea that some variables mediate the effect of preference reversals – there are two regression regimes (and that the

(known as E step) and then re-estimating the model parameters using the completions (known as the M-Step).

[49] For a formal exposition of the model see Appendix D.

[50] Also known as Root Mean Square Error.

mixture procedure is not simply 'creamskimming', or just putting the outliers in one group and those observations on or near the regression line in the other group).[51]

As for our main hypothesis test, the table shows that – while most other coefficients are remarkably similar across groups – preference reversals have a different effect in the two groups: whereas preference reversals are associated with no statistically significant increase in the seasonal difference in self-disconnections in group 1; they are associated with a significant increase (by 89%) in group2.

7.4 Second Step

In the last section(s) we showed that there are two distinct groups of households in our sample. In addition, we showed that preference reversals have a much stronger effect on households' ability to smooth consumption in the second group than in the first one.

We can now test our hypothesis that the reason for the difference between the two groups is the presence/absence of 'sophistication' and/or 'interference'. We do this by classifying households as either a member of group 1 or group 2 using the (posterior) probability from the first step[52] and then regressing group membership on a set of variables – including proxies for 'sophistication' and 'interference'.

The basic idea is that if the difference in parameter estimates in the last section are due to differences in 'sophistication' and/or 'interference', we would expect to find that individuals in group 1 are more 'sophisticated' and/or more likely to be subject to 'interference' than individuals in group 2.

Formally, we estimate:

$$GMS_i = \delta_0 + \delta_1 I_i + \delta_2 Z_i + \eta_i \quad (13)$$

where GMS_i is group membership of household i. I_i is a vector summarising our main variables (measuring 'sophistication' and 'interference'); Z_i a vector of control variables and δ the corresponding regression coefficients. η_i is an error term.

[51] More evidence in support of the existence of two regimes is provided in Appendix C.

[52] We will use different cut-off points (in probabilities) to ensure the robustness of our findings.

7.5 A Note on the data

Similar to Ameriks (2007), we measure the degree of sophistication (of the household member in charge of purchasing top-up) by presenting individuals with the following hypothetical scenario:

> "Suppose that you win 10 vouchers, each of which can be used for an evening out. On each such evening, you and a companion will get an unlimited budget for food and drinks at a place of your choosing. There will be no cost to you. The vouchers are available for immediate use, starting today and there is an absolute guarantee that they will be honoured by any place you select, if they are used within the next two years. However, if they are not used within this two year period, any vouchers that remain are valueless."

and then asking: how many of the vouchers would you ideally like to use in the first year? and: would you want to make this choice irreversible? The idea of the second question is that if a household member is aware of his/her preference reversal (i.e. is 'sophisticated'), he/she will want to tie his/her hands – to avoid that he/she departs from his/her long-term plan of action (once the time comes).

To measure the degree of 'interference', we use an approach suggested by Brutscher (2011). The idea is that in order for a partner to be able to influence the behaviour of the household member in charge of purchasing top-up, two conditions need to hold: i) the partner must display a stronger preference to smooth consumption/self-disconnection and ii) he/she must have a stronger bargaining position.

One way to capture the two conditions is by asking households 'who gives higher priority to saving for heating in the autumn/winter?' and 'for expenses other than heating who decides on when to buy expensive things?'. We then capture the effect of a household satisfying both conditions by using an interaction term between the two variables (which takes a value of '1' if the partner 'gives more priority to saving' and also 'tends to decide on when to buy expensive things' and '0' otherwise).[53]

[53] As discussed in Brutscher (2011), if the household member in charge of purchasing top-up is aware of his/her preference reversals, it may not be necessary for his/her partner to have a stronger bargaining position than him/her (to exercise 'interference'). To test this possibility, we redid the entire following analysis focusing only on 'differences in preferences'. We arrived at the same conclusions.

7.6 Analysis

Table 9 below provides the results from our estimation of (8).

	Pr[Group 1]$\geq$ 0.5	Pr[Group 2]$\geq$ 0.66	Pr[Group 2]$\geq$ 0.75
Hyperbolic	0.04*	0.07***	0.05**
	(0.03)	(0.03)	(0.02)
Interaction	0.02	0.10	0.05
	(0.09)	(0.11)	(0.10)
Higher Priority	-0.01	0.09	-0.06
	(0.03)	(0.09)	(0.08)
Who Decides	-0.04	-0.05	-0.04
	(0.05)	(0.03)	(0.03)
Socio-economic Controls	Yes[54]	Yes	Yes

Table 9: Regression Analysis of (8) - Statistically Significant at 1% ***; 5% **; 10% *

Column (1) shows the results if we classify households with a (posterior) probability of ≥ 0.5 as members of group 1; column (2) shows the results if we classify households with a (posterior) probability ≥ 0.66 as members of group1; and column (3) the results if we classify households with a (posterior) probability ≥ 0.75 as members of group1.

What we find is that – independent of how we classify households – households in the first group are significantly more likely to be 'sophisticated' than households in the second one – which suggests that being aware of one's own preference reversals is a big step towards addressing the problem.[55]

In contrast, households in the first groups are no more likely to satisfy the conditions of 'interference'. This suggests that household members not suffering from preference reversals are not as good in mediating the negative consequences of their partner' preference reversals (and that policy makers may have to provide them with a commitment savings device).

[54] None of the control variables are statistically significant - suggesting that 'sophistication' is the only mediating factor in our dataset.

[55] An alternative explanation for the presence of two regimes (in Table 8) and our findings in Table 9 is that some households live further away from the nearest outlet selling top-up than others. To the extent that distance matters more in the autumn/winter (when the weather tends to be bad) than the spring/summer, this may affect our findings in Tables 8 and 9. It is important to note, however, that we control for 'distance from the nearest outlet' in both stages (FMM and analysis of posterior) - and in neither case find a significant effect.

One possible explanation for this finding is that household members who are 'sophisticated' are more likely to have accepted their 'weakness' – and so need a less strict mechanism/commitment device to 'lock up' their savings – than household members which are (only) subject to 'interference'.

8 Conclusion and Policy Implications

In this paper, we revisited the problem of self-disconnection among electricity pre-payment customers in Great Britain.

In the first half of the paper we discussed the extent and drivers of self-disconnection: using metering data from 2.3 million pre-payment households from British Gas and data from a telephone survey of 502 (randomly selected) households with a pre-payment meter, we showed that:

1. The majority of households with a pre-payment meter do not self-disconnect (78%); 12% of households tend to self-disconnect once; approx. 3% of households tend to self-disconnect more often than 4 times.

2. Proxying the duration of self-disconnections, we also found that most self-disconnections (>62%) last for less than one day; between 72% and 82% of self-disconnections last for less than two days.

3. Finally, we showed that the main driver of self-disconnections is 'financial constraints' – which suggests that reducing the total number of self-disconnections is likely to be difficult/expensive.

In the second half of the paper, we discussed an alternative/complementary approach to reducing the total number of self-disconnections: we explored whether it is possible to improve the timing of households' self-disconnection behaviour.

We proceeded in three steps:

1. In the first step, we argued that households' self-disconnection behaviour over time tends to be sub-optimal.

2. In the second step, we provided evidence which suggests that one important driver for household's sub-optimal self-disconnection behaviour is (hyperbolic) preference reversals.

3. In the final step, we discussed several possible policy responses to preference reversals (in the context of smoothing self-disconnections over time) and concluded that an effective policy is likely to include an initiative to increase awareness of preference reversals and possibly the provision of a commitment savings device.

Taking this work forward, the main tasks will include: i) designing an awareness campaign; ii) a commitment savings device and iii) a research design which allows to evaluate the effectiveness of these policies with regard to households' ability to smooth self-disconnections over the course of a year/with regard to reducing the negative effect of self-disconnections.

9 Bibliography

References

[1] Abadie, A. (2003). Semiparametric instrumental variable estimation of treatment response models. Journal of Econometrics, 113(2), 231–263. Elsevier. Retrieved from http://www.sciencedirect.com/science/article/pii/S0304407602002014

[2] Agarwal, S., & Driscoll, J. (2008). Learning in the credit card market. Retrieved from http://www.nber.org/papers/w13822

[3] Akerlof, G. (1991). Procrastination and obedience. The American Economic Review. Retrieved from http://www.jstor.org/stable/2006817

[4] Angrist, J. D., & Pischke, J. S. (2008). Mostly harmless econometrics: An empiricist's companion. Princeton Univ Pr. Retrieved from http://books.google.co.uk/books?hl=en&lr=&id=ztXL21Xd8v8C& oi=fnd&pg=PR7&ots=Uh36Tx-6JuW& sig=eA7V7wysc7NP8XLwspFtrA7oArQ

[5] Ariely, D, Wertenbroch, K. (2002). Procrastination, deadlines, and performance: Self-control by precommitment. Psychological Science. Retrieved from http://pss.sagepub.com/content/ 13/3/219.short

[6] Ashraf, N. (2006). Deposit collectors. The BE Journal of Economic Analysis & .... Retrieved from http://www.bepress.com/bejeap/advances/vol6/iss2/art5/

[7] Basu, K. (2009). Seasonal Food Consumption in West Timor! Retrieved from http://real.wharton.upenn.edu/~maisy/documents/ Basu-Wong_NEUDC_Aug2009_Final.pdf

[8] Benzion, U. (1989). Discount rates inferred from decisions: An experimental study. Management science. Retrieved from http://www.jstor.org/stable/2631972

[9] Birmingham Settlement, Community Energy Research and Bristol Energy Centre (1993) An Investigation of Consumer, Fuel Company and Agency Responses to Prepayment Meters and Fuel Direct. Birmingham Settlement.

[10] Brocas, I., & Carrillo, J. D. (2004). The psychology of economic decisions, Volume 2 (p. 231). Oxford University Press. Retrieved from http://books.google.com/books?hl=en&lr=&id=2QlfifCIhtsC&pgis=1

[11] Brutscher, P. B. (2011). Making Sense of Oil Stamp Savings Programmes. EPRG Working Paper Series.

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

[13] Commissaris, C. (1998). Subjective forgetfulness in a normal Dutch population: possibilities for health education and other interventions. Patient education and counseling. Retrieved from http://www.sciencedirect.com/science/article/pii/S0738399198000408

[14] Cooke, D., Ferrari, A., & Giulietti, M. (2001). Affording Gas and Electricity: Self disconnection and rationing by prepayment and low income credit consumers and Company attitudes to social action. Electricity Association. Retrieved from http://scholar.google.de

[15] Deb, P., & Trivedi, P. K. (1997). Demand for medical care by the elderly: a finite mixture approach. Journal of applied Econometrics, 12(3), 313–336. Wiley Online Library. Retrieved from http://onlinelibrary.wiley.com/doi/10.1002/(SICI)1099-1255(199705)12:3%3C313::AID-JAE440%3E3.0.CO;2-G/abstract

[16] DellaVigna, S; Malmendier, U. (n.d.). Paying not to go to the gym. Citeseer. Retrieved from http://citeseerx.ist.psu.edu/viewdoc/download ?doi=10.1.1.61.4827&rep=rep1&type=pdf

[17] Doble, M. (2000). A regulatory policy for self-disconnection: an examination of the reasons for and implications of pre-payment meter stoppages. Policy Studies. Retrieved from http://www.ingentaconnect.com/content/routledg/cpos/2000/00000021/00000003/art00004

[18] Drakeford, M. (1997). The poverty of privatization. Critical Social Policy. Retrieved from http://csp.sagepub.com/content/17/51/115.short

[19] Federation, N. H. (2008). Pre-Payment Meter Utilities Customers.

[20] Geweke, J., & Petrella, L. (1998). Prior Density-Ratio Class Robustness in Econometrics. Journal of Business & Economic Statistics, 469–478. JSTOR. Retrieved from http://www.jstor.org/stable/ 1392615

[21] Gibbons, D., & Singler, R. (2008). Cold Comfort: A review of coping strategies employed by households in fuel poverty. Report for Energywatch, Inclusion Research and Consultancy. Retrieved from http://scholar.google.de/scholar?q=Cold+Comfort%3A++A+Review+of +Coping++Strategies+Employed+by+Households+in+Fuel++ Poverty&hl=de&btnG=Suche&lr=#0

[22] Gritz, R. M. (1993). The impact of training on the frequency and duration of employment. Journal of Econometrics, 57(1-3), 21–51. Elsevier. Retrieved from http://www.sciencedirect.com/science/ article/pii/030440769390057C

[23] Hausman, J. (1984). Econometric models for count data with an application to the patents-R&D relationship. Retrieved from http://www.nber.org/papers/t0017

[24] Heckman, J. J., Robb, R., & Walker, J. R. (1990). Testing the mixture of exponentials hypothesis and estimating the mixing distribution by the methods of moments. Journal of the American Statistical Association, 582–589. JSTOR. Retrieved from http://www.jstor.org/stable/2289802

[25] Koenker, R.; Bassett, G. (1978). Regression quantiles. Econometrica: journal of the Econometric Society. Retrieved from http://www.jstor.org/stable/1913643

[26] Koenker, R.; H. K. (2001). Quantile regression. The Journal of Economic Perspectives. Retrieved from http://www.jstor.org/stable/2696522

[27] Laibson, D. (1997). Golden Eggs and Hyperbolic Discounting*. Quarterly Journal of Economics. Retrieved from http://www.mitpressjournals.org/doi/pdf/10.1162/003355397555253

[28] Lee, S.(2007). Endogeneity in Quantile Regression models: A Control Function Approach. Journal of Econometrics; 141(2), 1131-1158

[29] Machado, J; Silva, S. (2005). Quantiles for counts. Journal of the American Statistical Association. Retrieved from http://pubs.amstat.org/ doi/abs/10.1198/016214505000000330

[30] McLachlan, G. (2000). Finite mixture models. Retrieved from http://books.google.de/books?hl=de&lr=&id=c2_fAox0DQoC&oi=fnd&pg=PR19&dq=McLachlan,+G.J.,+and+D.+Peel.+2000.+%E2%80%9CFinite+Mixture+Models&ots=IrOFZX52kA&sig=w6vmXu3Ld9mOIlQld1MamHS-5G0

[31] Merseyside Right to Fuel Action Group and Liverpool City Council (1998) The Liverpool Fuel Poverty Survey. Liverpool City Council Anti-Poverty Team.

[32] Morduch, J. J., & Stern, H. S. (1997). Using mixture models to detect sex bias in health outcomes in Bangladesh• 1. Journal of Econometrics, 77(1), 259–276. Elsevier. Retrieved from http://www.sciencedirect.com/science/article/pii/S0304407696018155

[33] Mullahy, J. (1997). Heterogeneity, excess zeros, and the structure of count data models. Journal of Applied Econometrics. Retrieved from http://onlinelibrary.wiley.com/doi/10.1002/(SICI)1099-1255(199705)12:3%3C337::AID-JAE438%3E3.0.CO;2-G/abstract

[34] O'Donoghue, T; Rabin, M. (2001). Choice and Procrastination*. Quarterly Journal of Economics. Retrieved from http://www.mitpressjournals.org/doi/pdf/10.1162/003355301556365

[35] Paxson, C. (1993). Consumption and income seasonality in Thailand. Journal of Political Economy. Retrieved from http://www.jstor.org/stable/2138673

[36] Regulation, C. for M. under. (2001). Affording Gas and Electricity: Self Disconnection and Rationing by Prepayment and Low Income Credit Consumers and Company Attitudes to Social Action.

[37] Rowlingson, K.; Kempson, E. (1993). Gas debt and disconnections. Retrieved from http://opensigle.inist.fr/handle/10068/480977

[38] Scharff, R. (2009). Obesity and hyperbolic discounting: Evidence and implications. Journal of Consumer Policy. Retrieved from http://www.springerlink.com/index/96R18V46WK2H3765.pdf

[39] Scottish-Hydro-Electric (1997): Customer research. Card Meter Self Disconnection Survey. Postal Questionnaire. Loyalty Research and Development.

[40] Scottish-Hydro-Electric (1998) Card Meter Self Disconnection Survey. Customer Research Team, Energy Marketing.

[41] Questionnaire. Loyalty Research and Development.

[42] Titterington, D. M., A. F. M. S. and U. E. M. (1995). Statistical Analysis of Finite Mixture Distributions. New York: John Wiley.

[43] Tversky, A.; Kahneman, D. (1986). Rational choice and the framing of decisions. Journal of business. Retrieved from http://www.jstor.org/stable/2352759

[44] Wang, P., I.M. Cockburn, and M. L. P. (1998). Analysis of Patent Data - A Mixed Poisson Regression Model Approach. Journal of Business and Economic Statistics, 16(27-36).

[45] Wedel, M., DeSarbo, W. S., Bult, J. R., & Ramaswamy, V. (1993). A latent class Poisson regression model for heterogeneous count data. Journal of Applied Econometrics, 8(4), 397–411. Wiley Online Library. Retrieved from http://onlinelibrary.wiley.com/doi/10.1002/jae.3950080407/abstract

[46] Wolfe, J. H. (1971). A Monte Carlo Study of the Sampling Distribution of the Likelihood Ratio for Mixtures of Multinormal Distributions. Retrieved from http://oai.dtic.mil/oai/oai?verb=getRecord&metadataPrefix=html&identifier=AD0731039

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

10 Appendix A

Table 10 below provides summary statistics for our survey – by strata.

Variable	Strata 1	Strata 2	Strata 3	Strata 4
N	126	125	126	125
Age	45.65	43.68	41.67	43.61
	(13.87)	(13.51)	(12.06)	(12.31)
Gender	0.67	0.49	0.66	0.67
(Female=1)	(0.47)	(0.50)	(0.48)	(0.47)
Marital Status	0.34	0.46	0.39	0.35
(Single=1)	(0.48)	(0.50)	(0.49)	(0.48)
Income	161	181	219	187
(Weekly)	(158.81)	(157.42)	(166.01)	(148.86)
Labour Market Status	0.47	0.54	0.54	0.44
(Active=1)	(0.50)	(0.50)	(0.50)	(0.50)
Household Size	2.73	2.50	2.56	2.94
	(1.43)	(1.71)	(1.42)	(1.61)
Number of Children	0.90	0.82	0.90	1.09
	(1.18)	(1.24)	(1.06)	(1.30)
Type of House	0.06	0.11	0.08	0.12
(detached=1)	(0.24)	(0.32)	(0.27)	(0.33)
Rent	0.74	0.81	0.70	0.75
(Rent=1)	(0.44)	(0.40)	(0.46)	(0.44)

Table 10: Summary Statistics - by Strata.

11 Appendix B

In section (5.5) we showed that the seasonal pattern in self-disconnection is more pronounced among households that face a strong seasonal pattern in income/energy use than among households that do not face such a strong pattern in income/energy use.

Here we provide the corresponding quantiatitve estimates (and F-tests). The underlying regression models take the following form:

$$A_{imt} = \alpha_m + \alpha_m^F GPPM_{it} + \alpha GPPM_{it} + \gamma X_{it} + \varepsilon_{imt} \quad (14)$$

for energy expenditure (A_{imt}) of household i in month m of year t - where the coefficients α_m capture month effects for households paying for their gas by means of a budgeting scheme and α_m^F the additional month effects for households which have a gas pre-payment meter; α captures the difference in the base level of expenditure - and

$$SD_{imt} = \beta_m + \beta_m^F Z_{it} + \beta Z_{it} + \phi_i + u_{imt} \quad (15)$$

for self-disconnection (SD_{imt}) of household i in month m of year t - where the coefficient β_m captures month effects for households paying for their gas by means of a budgeting scheme; and β_m^F captures the additional month effects for households with a gas pre-payment meter.

To allow for different (base) levels of self-disconnections between the two groups of households, we include a dummy variable equal to one if a household is exposed to seasonal changes in income/ energy use and zero otherwise (again). β captures the corresponding difference in (base) levels. Finally, we include individual fixed effects (ϕ_i) to control for time-invariant differences across individuals.

Table 11 provides the results from estimating (14) and (15):

	Income/Energy Use		Self-Disconnection	
	Month	Month*PPM$_{gas}$	Month	Month*PPM$_{gas}$
Jan	16.18	-8.98	-0.0057*	0.0023
	(14.79	16.14	(0.0033)	(0.0034)
Feb	6.17	-7.85	-0.0064*	0.0033
	(15.96)	(17.21)	(0.0033)	(0.0034)
Mar	2.03	-1.51	0.0036	-0.0028
	(15.95)	(17.28)	(0.0033)	(0.0034)
Apr	0.15	-19.30	-0.0059*	0.0004
	(15.33)	(16.90)	(0.0033)	(0.0034)
May	5.42	-28.22*	-0.0025	-0.0049*
	(14.82)	(16.29)	(0.0033)	(0.0034)
June	14.66	-41.38**	0.0020	-0.0064*
	(17.96)	(19.23)	(0.0033)	(0.0034)
July	-5.13	-30.66**	0.0010	-0.0074**
	14.08	(15.76)	(0.0033)	(0.0034)
Aug	2.45	-35.51**	0.0017	-0.0073**
	(14.10)	(15.67)	(0.0033)	(0.0034)
Sep	13.02	-40.72***	0.0005	-0.0027
	13.94	(15.42)	(0.0033)	(0.0034)
Oct	-6.96	-6.38	-0.0028	0.0017
	(15.36)	16.61	(0.0033)	(0.0034)
Nov	12.68	16.61	0.0005	0.0042
	(13.90)	(15.36)	(0.0033)	(0.0034)

Table 11: Quantative Estimates of FNX and FNX -
Statsitically Singicant at *** 1%; ** 5%; * 10%

The numbers under the column 'month' are the month effects for households with no pre-payment meter. The numbers under the column '$month*PPM$' are the additional month effects for households which do have a gas pre-payment meter – and therefore represent the difference in monthly income/energy use and self-disconnections between the two groups of households.

Table 12 provides a series of F-tests to test the statistical significance of the differences in seasonal patterns in the two groups.

Dependent Variable	Test	
Income/Energy Use	Test 1: no month effect - gas	F=9.05
	pre-payment meter	p<0.01
	Test 2: no month effect - gas	F=0.61
	budgeting scheme	p=0.82
	Test 3: month effects	F=1.57
	identical	p=0.10
Self-Disconnections	Test 1: no month effect - gas	F=43.75
	pre-payment meter	p<0.01
	Test 2: no month effect - gas	F=2.39
	budgeting scheme	p=0.01
	Test 3: month effects	F=3.00
	identical	p<0.01

Table 12: F-Tests of Seasonal Patterns

Tests 1 and 4 in Table 13 show that the month effects ($\alpha_m + \alpha_m^F$ and $\beta_m + \beta_m^F$) in our income and self-disconnection equation are jointly significant – and so that for households with a gas pre-payment meter both the seasonal patterns in income/energy use and the seasonal pattern in the probability to self-disconnect are (not only visible but also) statistically significant.

Test 2 shows that – for households without a gas pre-payment meter – the month effects (α_m) are jointly insignificant (which is in line with our hypothesis that there is no seasonal pattern in income/energy use). When it comes to the probability to self-disconnect, on the other hand, we cannot exclude that (also) households without a gas pre-payment meter display a seasonal pattern in the probability to self-disconnect (Test 3).

What is important to note, however, is that despite the fact that we find a seasonal pattern in the probability to self-disconnect for households with a gas pre-payment meter and households with no gas pre-payment meter, we find strong evidence that the two patterns are not identical (Test 6). Instead, the pattern for households with a gas pre-payment meter is significantly more pronounced than the one for households without a gas pre-payment meter.

12 Appendix C

In this section, we discuss several tests of the identification assumption in section (5.5).

12.1 Robustness Test 1: Identification Assumption prices and/or tastes

One of the assumptions we tested in section 5.5 is that the seasonal pattern in self-disconnections is not driven by seasonal differences in the duration of self-disconnections.

If self-disconnections last longer in the spring/summer than the autumn/winter, it is possible that we find a strong seasonal pattern in the probability that households self-disconnect in a particular month – even if households do (in fact) smooth consumption/the total amount of self-disconnection over the course of a year.

In this section, we provide complementary evidence to the findings in section (5.5) by showing that there are (indeed) no seasonal differences in the duration of self-disconnections: in Figure 9a and 9b below, we plot the distribution of when households top-up their meters after (being predicted to) having self-disconnected – separately for spring/summer and autumn/winter.

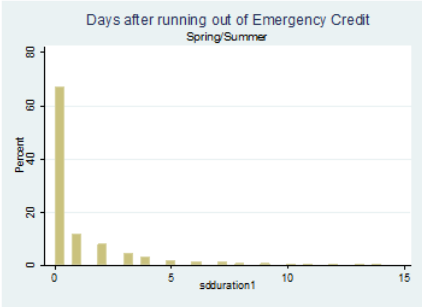

Figure 9a: Duration of Self-Disconnection – Spring/Summer

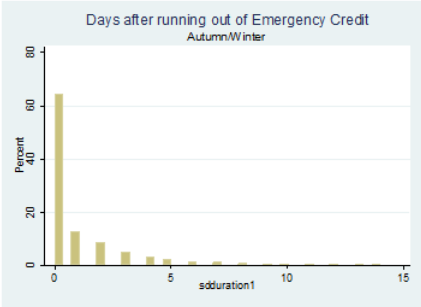

Figure 9b: Duration of Self-Disconnection – Autumn/Winter

What we find is an almost identical pattern in the two seasons: The figures show that – in both seasons – the large majority of self-disconnections last for less than 1 day (with households topping-up their meter on the same day they are predicted to have run out of credit); between 72% and 82% of self-disconnections last less than 2 days.

The similarity between the two plots is at odds with the idea that the seasonal pattern in self-disconnections (and, hence, the comovement between the probability that households self-disconnect and their income flows/energy use) is due to differences in the duration of self-disconnections across seasons.

12.2 Robustness Test 2: Identification Assumption prices and/or tastes

The second key identification assumption underlying our analysis in section (5.5) is that prices and/or seasonal preferences are unrelated to whether households have a gas pre-payment meter or not.

If households with a gas pre-payment meter face a stronger increase in the price for credit/to avoid self-disconnection in the autumn/winter than households on a budgeting scheme, then the stronger increase in the number of self-disconnections (for this group) is entirely consistent with the idea that the seasonal pattern in self-disconnections is due to seasonal differences in prices.

Similarly, if households with a gas pre-payment meter have a stronger preference to self-disconnect in the autumn/winter than households on a budgeting scheme, then the stronger increase in the number of self-disconnections among households with a gas pre-payment meter is entirely consistent with the idea that the seasonal pattern in self-disconnections is due to differences in tastes.

Figures 10a and 10b show, however, that both possibilities are at odds with the data: Figure 10a plots the (average)[56] gas tariff for households with a gas pre-payment meter and households on a budgeting scheme from 2007 to 2010.[57] The figure shows that the tariffs for the two groups are highly correlated – and provides no evidence for a stronger increase in tariffs in the autumn/winter for households with a gas pre-payment meter.

[56] Taking into account geographic differences in tariffs.
[57] To facilitate comparison, we normalise 2007 tariffs for both groups to 100.

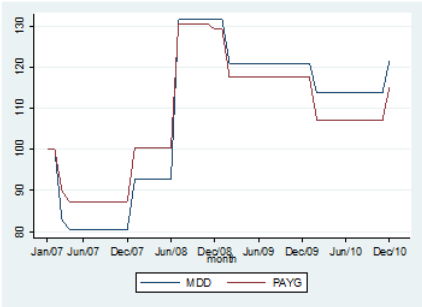

Figure 10a: Tariff – Gas PPM cust. and Gas BS customers

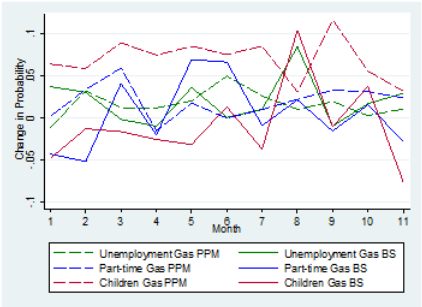

Figure 10b: Probabs. – Gas PPM cust. and Gas BS customers

Figure 10b plots the probability i) that a household head in our sample is unemployed; ii) that he/she is employed part-time; and iii) that one or more children live in his/her household – separately for households with a gas pre-payment meter and households on a budgeting scheme.

The figure shows that for all three variables the patterns (in probabilities) are similar for households with a gas pre-payment meter and households on a budgeting scheme: that is, we find no evidence that households with a gas pre-payment meter show a stronger increase in the probability to be unemployed; working part-time; or to have children in the autumn/winter than households on a budgeting scheme.

This suggests that it is rather unlikely that these households have a stronger preference/taste to self-disconnect in the autumn/winter – because they tend to spend less time at home during the autumn/winter. Similarly, it suggests that it is unlikely that households with a gas pre-payment meter have a stronger preference/taste to self-disconnect in the autumn/winter – because they are less vulnerable than households on a budgeting scheme.

12.3 Robustness Test 3: Identification Assumption Differences in the Number of Self-Disconnections

Another (implicit) assumption underlying our analysis in section (5.5) is that households on a budgeting scheme do self-disconnect.

If households on a budgeting scheme do not self-disconnect/self-disconnect rarely (e.g. because they are better off than households with a gas pre-payment meter) this implies – independently from the true driver of the seasonal pattern in the probability to self-disconnect – a flat(ter) pattern in the probability that

households self-disconnect.

To test this possibility, we re-run our analysis (from sections 5.5) – this time using only households which have self-disconnected at least four times over the course of the past 4 years. The findings from this analysis are presented in Figure 11a and Figure 11b below.

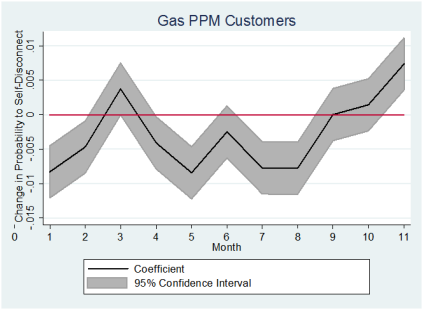

Figure 11a: SD+4 -Gas-PPM Customers

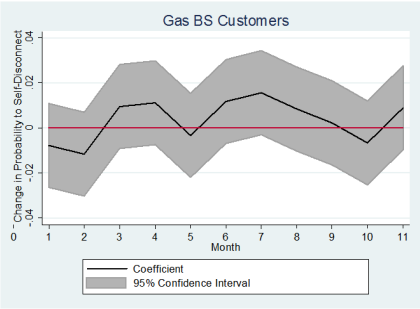

Figure 11b: SD+4 -Gas Budgeting Scheme

What we find is the same, flatter pattern in the probability that households self-disconnect (over the course of a year) among households on a budgeting scheme than households with a gas pre-payment meter as before.

13 Appendix D

In this section, we give a formal exposition of our finite mixture model and provide several specification tests.

13.1 Specification Finite Mixture

The incomplete, or observed, data likelihood of our model is then given by:[58]

$$f(Y_i) = \frac{\theta}{\sqrt{2\pi\sigma_1^2}} \exp\left\{-\frac{(Y_i - X_i\beta_1)^2}{2\sigma_1^2}\right\} + \frac{1-\theta}{\sqrt{2\pi\sigma_2^2}} \exp\left\{-\frac{(Y_i - X_i\beta 2)^2}{2\sigma_2^2}\right\} \quad (10)$$

To write the complete-data likelihood, we have to define the indicator variable d_{ij} where d_{i1}=1 if the observation is associated with the first group, 0 otherwise, and d_{i2}=1 (in our two-group case, really $1 - d_{i1}$=1) if the observation is associated with the second group, 0 otherwise.

The problem is that – as discussed – d is not/or only imperfectly observed. This means that d must be considered a random variable and that the complete-data density function is given by:

$$f(Y_i) = \left\{\frac{\theta}{\sqrt{2\pi\sigma_1^2}} \exp\left\{-\frac{(Y_i - X_i\beta_1)^2}{2\sigma_1^2}\right\}\right\}^{d_{1i}} \left\{\frac{1-\theta}{\sqrt{2\pi\sigma_2^2}} \exp\left\{-\frac{(Y_i - X_i\beta 2)^2}{2\sigma_2^2}\right\}\right\}^{1-d_{1i}} \quad (11)$$

To estimate (11), we can use a two-step (iterative) procedure called the Expectation Maximisation Algorithm: in the first step, we replace d by its expectation given the data. This expectation is given by $E(d_{i1}|Y_i) = (1)[P(d_{i1} = 0|Y_i)] + (0)[P(d_{i1} = 0|Y_i)]P(d_{i1} = 1|Y_i)$.

The expected value can then be evaluated using Bayes' rule:

$$P(d_{1i} = 1|Y_i) = \frac{P(d_{1i}=1)P(Y_{1i}|d_{1i}=1)}{\sum_{i=1}^{2}(P_{ij}=1)P(Y_i|d_{ij}=1)} = \frac{\theta f_{i1}}{\theta f_{i1} + (1-\theta) f_{i2}} = w_{i1} \quad (12)$$

In the second step, we substitute the expected values/weights from the first step into the log of our complete-data likelihood and maximise it with respect to the unknown parameters. We then use the resulting estimates to update the expectations in (12), and hence the probability that a household belongs to one of the two groups of households – which gives us a new set of estimates of the unknown parameters.

[58] These densities comprise the logarithm of the complete-data likelihood function that is given by: $\ln L = \sum_{i=1}^{n} \left\{d_{i1}(\ln\theta - \ln f_{i1}) + (1 - d_{i1})(\ln(1-\theta) + \ln f_{i2})\right\}$ where f_{i1} and f_{i2} are the respective normal density functions.

We repeat this procedure until no improvement in fit of our model (by further updating the probability/weight estimates and/or parameter estimates) is possible.[59]

13.2 Number of components

The first set of checks (re-)examines the number of components. We initially investigate a FMM specification with two heterogeneous groups. Then we test to determine whether one, two or three components better represent the underlying relationships. The test criteria for each model are reported in Table 13 while the resulting FMM estimates of being hyperbolic can be seen in Table 14.

We use several criteria in selecting the best model: AIC/BIC and appropriate mechanism of the estimation approach. The AIC/BIC reveal that the one component regression is inferior to a specification with multiple components. This is consistent with our finding that there are two groups with importantly different effects with regard to being hyperbolic. As a comparison, we allow three groups in the FMM specification. The three component model, however, does not converge.

	AIC	BIC
One Component Model	375.7218	397.5478
2 Component Model	375.35	382.630
3 Component Model	not converging	

Table 13: Selection Criteria for various models.

The coefficients of our measure of hyperbolic preference reversals (for the different models) is shown in Table 14 below:

[59] As with most optimisation methods for non-concave functions, the expectation maximisation algorithm comes with guarantees only of convergence to a local maximum of the objective function (except in degenerate cases). We deal with this problem by running the procedure using multiple initial starting parameters and by initialising parameters in a way that breaks symmetry.

	Overall	Component 1	Component 2	Component 3
One Component Model	0.39**	X.XX	X.XX	X.XX
2 Component Model	X.XX	0.25	0.89***	X.XX
3 Component Mode	not converging			

Table 14: Comparison of the effect of hyperbolic preference reversals across models - Statistically Significant at 1% ***; 5% **; 10% *

What we find is that – in line with our hypothesis – the first group is significantly more likely to ...

14 Appendix E

One possible problem with our analyses in section 6.4 and 7.3. is that – (hyperbolic) preference reversals may reflect something else than temporary short-sightedness. Four alternative explanations include: i) noise; ii) inability to understand the questions; iii) the timing of cash flows; and iv) transaction costs/lack of trust.

The first alternative explanation for hyperbolic preference reversals is noise: if households are not sure about their preferences or report their preferences with some error, it is possible that the reason why certain households appear more impatient now than later is chance. There are, however, at least two reasons why this is unlikely:

First of all, if people are not sure about their preferences, it is not clear why almost twice as many households should report time-inconsistencies in the hyperbolic direction than in the other direction. Secondly, if the preference reversals we find were due to noise only, then they should not predict real behaviour, such as whether households self-disconnect more in the autumn/winter or not.

A second possible explanation for the (hyperbolic) preference reversals we find is the inability to understand our hypothetical questions. If less-educated individuals are more likely to report preference reversals (in either direction) and less-educated individuals are more likely to self-disconnect primarily in the autumn/winter, then we would spuriously conclude that the seasonality in self-disconnections is due to hyperbolic preferences, rather than just a lack of education.

The problem with this argument is – again – that it is not clear why an inability to understand our hypothetical question should lead almost twice as many households to report preference reversals in hyperbolic direction than in the other direction. In addition, if an inability to understand the hypothetical questions were at the core of our finding, we would expect a positive correlation between whether someone is time-inconsistent and their level of education.

This is not the case: Table 15 Columns (1) and (2) below shows the results from a simple logit model linking whether or not a household is time-inconsistent with their level of education.

	Coefficient	Sim. Prob.	Coefficient	Sim. Prob
College or more	-0.57	-0.11	-0.56	-0.11
	(0.36)	(0.07)	(0.36)	(0.07)
Number of hh members	-0.01	-0.01	-0.01	-0.01
	(0.07)	(0.05)	(0.07)	(0.05)
HH income	0.01	0.007	0.01	0.01
	(0.05)	(0.04)	(0.05)	(0.04)
Age	0.002	0.01	0.002	0.008
	(0.009)	(0.03)	(0.01)	(0.04)
Low income now, high	XX	XX	-0.06	-0.01
income in 6 m			(0.23)	(0.05)
High income now, low	XX	XX	0.12	0.03
income in 6 m			(0.38)	(0.09)
Female	-0.08	-0.02	-0.08	-0.02
	(0.21)	(0.04)	(0.21)	(0.05)
Marital Status	-0.21	-0.04	-0.20	-0.04
(Single=1)	(0.24)	(0.05)	(0.24)	(0.05)

Table 15: Robustness Check of our Measure of Preference Reversals

The table shows that have a college degree or higher does not affect (positively or negatively) whether respondents display preference reversals. (The corresponding regression coefficient is statistically insignficant).

A third possible explanation for the (hyperbolic) preference reversals is the timing of cash flows: households which report patience (impatience) now and impatience (patience) later may be flushed with cash now (later) but expect to be short of cash later (now). In order to make sense, such a story also requires some element of savings (borrowing) constraint.

To test this possibility, we asked households what months are their high-income and low-income months. We then linked whether they reported time-inconsistencies (in either direction) to this response to this question. Our findings are summarised in Table 15 Columns (3) and (4). What we find is that expecting a higher (lower) income in 6 months has no significant effect on whether respondents display preference reversals.

Finally, it is possible that the (hyperbolic) preference reversals in our data is lack of trust/transaction costs. For instance, Fernandez-Villaverde and Mukherji (2002) argue that uncertainty in future rewards will lead individuals to choose

immediate rewards. This seems unlikely, however, since such an explanation – again – cannot explain why (hyperbolic) preference reversals are correlated with the seasonal pattern in self-disconnections.

To summarise, although it is possible that our measure of (hyperbolic) preference reversals reflects something else but temporary shortsightedness, we find very little evidence to be suspicious that it does.