

Minerva Access is the Institutional Repository of The University of Melbourne

Author/s:

Chan, Edwin Sing-Long

Title:

Essays on energy demand and behaviour

Date:

2021

Persistent Link:

<https://hdl.handle.net/11343/278747>

Terms and Conditions:

Terms and Conditions: Copyright in works deposited in Minerva Access is retained by the copyright owner. The work may not be altered without permission from the copyright owner. Readers may only download, print and save electronic copies of whole works for their own personal non-commercial use. Any use that exceeds these limits requires permission from the copyright owner. Attribution is essential when quoting or paraphrasing from these works.

Essays on Energy Demand and Behaviour

by

Edwin Sing-Long Chan

ORCID: 0000-0002-5416-2326

A thesis submitted in total fulfillment for the
degree of Doctor of Philosophy - Business and Economics

in the

Department of Economics

Faculty of Business and Economics

THE UNIVERSITY OF MELBOURNE

June 2021

Abstract

This thesis explores consumers' behaviour in the residential electricity industry. It contains three essays that (i) provide empirical evidence that consumers exhibit inattention and behaviour that is consistent with non-standard preferences and (ii) examine whether behavioural interventions are effective to affect consumers' behaviours.

In essay one, I investigate whether consumers are attentive to their energy usage in the absence of real-time information feedback and whether attention varies over the billing cycle. In the residential electricity industry, consumers are billed periodically, so in the absence of feedbacks between bills, they have to be attentive to their usage within the billing cycle if they want to make fully informed consumption decisions. To assess consumers' attention, I ask them to predict their electricity bills in a survey. Since consumers pay for the cumulative sum of their usage within the billing cycle, they have a different amount of information (their incurred usage) to predict their bills depending on when they are in their billing cycles. I exploit the random variation in consumers' billing cycle starting date, so consumers have an exogenous variation in the amount of information to predict their bills, depending on when in the billing cycle a consumer responds to my survey. I expect consumers who respond to the survey at the end of the billing cycle to have an error-free prediction if they are perfectly attentive. I find that prediction errors fall for the first 12 days of the billing cycle, suggesting consumers are attentive to their current billing cycle's usage for 12 days. This is consistent with the previous findings, that receiving a previous bill in the early stage of the billing cycle makes electricity expenditure salient. However, this saliency does not seem to endure for the whole billing cycle.

Essay two investigates whether real-time information feedback in conjunction with goal-setting affects consumers' energy usage. Estimating this effect is challenging because the most common implementation of goal-setting is through voluntary participation. Hence, we cannot directly compare goal-setters and non-goal-setters due to selection. I partner with a local electricity retailer and implement a field experiment that randomly allocates consumers with an App that provides them with real-time feedback. In addition, I provide an option for consumers to set an energy usage goal. To address selection into goal-setting, I use a recently developed method in machine-learning, where I leverage the high-frequency feature of my smart meter energy usage data and predict the consumer's energy usage as a counter-factual. I find goal-setting in conjunction with real-time information feedback generates on average a \$12 annual saving for a consumer. Without carefully controlling for household-time varying confounders would overstate the effect of goal-setting with real-time information by 64%.

Essay three aims to understand the behavioural mechanism that drives the estimated effect from essay two. I propose and estimate a model of decision making where the decision-maker exhibits reference-dependent utility and costly information acquisition. The decision-maker chooses daily energy usage and information acquisition within a billing cycle, subject to a goal. I find that consumers on average reduce their energy consumption by 13% to avoid the disutility from missing their goals. While information acquisition cost plays a big role in reducing consumers' probability to acquire information, not being fully informed in real-time does not affect energy usage. I speculate this is due to sample selection, as my estimation sample consists of energy-conscious consumers who are informed of their energy usage in the absence of real-time information feedback.

Declaration of Authorship

I, Edwin Sing-Long Chan, declare that this thesis titled, 'Essays on Behaviour and Energy Demand' and the work presented in it are my own. I confirm that:

- The thesis comprises only my original work towards the Doctor of Philosophy except where indicated in the preface;
- due acknowledgement has been made in the text to all other material used; and
- the thesis is fewer than 100,000 words in length, exclusive of tables, maps, bibliographies and appendices as approved by the Research Higher Degrees Committee.

Signed:

Edwin Sing-Long Chan

Date: June 2021

Preface

This thesis contains original research in Chapters 2 through 4.

- Chapter 2 is based on the following working paper:
Chan, E (2021). Forgetting and forecasting energy bills. *Working Paper*
- Chapter 3 is based on the following working paper:
Chan, E (2021). Machine learning from households about setting an energy goal. *Working Paper*
- Chapter 4 is based on the following working paper:
Chan, E (2021). Goals, information friction, and energy usage. *Working Paper*

Chapters 2 through 4 have been supported by the Australian Research Council under Linkage Grant LP140100099. This research was undertaken using the LIEF HPC-GPGPU Facility hosted at the University of Melbourne. This Facility was established with the assistance of LIEF Grant LE170100200.

Acknowledgements

I acknowledge the financial support from the University of Melbourne's Faculty of Business and Economics Doctoral Program Scholarship.

Throughout the PhD, I have received a great deal of support.

I am grateful to my supervisors, David Byrne and Maria Recalde for their support and guidance throughout my Ph.D. Your insightful feedback pushed me to sharpen my thinking and brought my work to a higher level. Thank you for being critical of my work in a healthy way, and keep encouraging me to take "good risks". I have learned so much from both of you, not only as Economists but also as genuine human beings.

I am also thankful for the department members at The University of Melbourne, especially the constant support from Kalvinder Shields and John DeNew throughout this time.

To my fellow Ph.D. colleagues - Haikun Zhan, Silvia Griselda, Shashi Karunanethy, Sabrina Tabassum, Alan Weiss, Khac Xuan Vu, Charles Siriban, and Tao Sun, thank you for being so inspirational and helpful. I will never forget all the hours we spent in 305.

There are so many friends I would like to thank, especially Ken Quach, Jonathan Lim, Jack Chen, Reynold Chan, David Wu, Sam Cheung, Jonathan Leong, John Tang, and Heiman Cheung. Thank you for all your messages and prayers. Also, a big thanks to Hans Zhu, who assisted me with data extraction.

I also wish to thank my family - Eric, Ella, and Edward, for providing me with a studious environment for me to grow up in and encouraged me to ask questions from a very young age. I would like to thank Alice. Thank you for always being there for me.

All praise and glory to my Lord, Jesus Christ.

Contents

Abstract	i
Declaration of Authorship	iii
Preface	iv
Acknowledgements	v
List of Figures	ix
List of Tables	x
1 Introduction	1
2 Forgetting and Forecasting Energy Bills	4
2.1 Introduction	4
2.1.1 Related literature	6
2.2 A thought experiment	7
2.3 Data	8
2.3.1 Background and context	8
2.3.2 Survey design and data source	9
2.3.3 Sample selection	11
2.3.4 Summary statistics	12
2.4 Are consumer attentive to their electricity usage?	14
2.4.1 Graphical evidence	14
2.4.2 Econometrics specification	15
2.4.3 Threats to identification strategy	17
2.4.4 Implication of the dynamics of attention	20
2.4.5 Investigating the use of available information	23
2.4.6 Does automatic payment matter?	25
2.5 Conclusion	27
3 Machine learning from households about setting an energy goal	29
3.1 Introduction	29
3.1.1 Related literature	31
3.2 The Experiment	32

3.2.1	Background and context	32
3.2.2	The platform	33
3.2.3	Experimental design	34
3.3	Conceptual framework	35
3.4	Data	36
3.4.1	Data sources	36
3.4.2	Who sets a goal?	38
3.5	Panel fixed effects approach	44
3.5.1	Results - DiD	45
3.6	Machine Learning approach	47
3.6.1	The first step	48
3.6.2	The second step	50
3.6.3	Validity checks and diagnostics	51
3.6.4	Results - ML-DiD	55
3.7	Robustness checks	57
3.8	Policy implications	59
3.9	Conclusion	60
4	Goals, information friction, and energy usage	62
4.1	Introduction	62
4.1.1	Related literature	64
4.2	Data	66
4.2.1	Reference dependence and rational inattention	66
4.2.2	Within-billing cycle dynamics	67
4.2.3	Effect of Email reminders	69
4.3	Model	70
4.3.1	Model setup	70
4.3.2	Solving the model	73
4.4	Numerical solution	75
4.5	Identification	78
4.6	Estimation	80
4.6.1	Other estimation details	81
4.7	Results	83
4.7.1	Parameters interpretation	83
4.7.2	Policy counterfactuals	86
4.8	Conclusion	86
5	Conclusion	89
A	Appendices for Chapter 2	92
A.1	Incentive compatibility from the pilot survey	92
A.2	Survey email and questionnaire	93
A.3	Balance test by survey waves	95
A.4	Derivation and inferences of turning points	98
A.4.1	Quadratic specification	98
A.4.2	Cubic specification	98

A.4.3	Quartic specification	99
A.5	Robustness checks	99
A.6	Testing differences in turning points for two subgroups	99
A.6.1	Quadratic specification	99
A.6.2	Cubic specification	100
B	Appendices for Chapter 3	101
B.1	App usage overtime	101
B.2	Geographic locations of our sample and weather stations	104
B.3	Regression results for diff-in-diff models by 3 hours block	104
B.4	Regression results for ML diff-in-diff models by 3 hours block	105
B.5	Robustness checks for ML diff-in-diff models by 3 hours block	105
C	Appendices for Chapter 4	106
C.1	Variables descriptions for the model	106
C.2	Modelling present-bias	107
C.3	Monte Carlo study	108
C.4	Loss aversion and information acquisition	109
	 Bibliography	 111

List of Figures

2.1	A sample of an electricity bill	10
2.2	Variation in the day of the billing cycle when the consumer responds to our survey.	14
2.3	Average RMSE prediction within days	15
2.4	What if consumers are attentive for the whole cycle?	21
2.5	Average RMSE for recall by different days	23
2.6	Average RMSE within different days by payment method	25
3.1	Sample of the platform	33
3.2	Graphic representation of the experimental design	34
3.3	An example of a Weather Underground weather station data web-page . .	37
3.4	Parallel trend assumption - in three hours blocks	45
3.5	3-hours-blocks breakdown, Left: β_1 , Right: β_2	47
3.6	Graphical illustration of identification (ideally)	50
3.7	Graphical illustration of identification (in practice)	50
3.8	Validity check - number of variables selected	53
3.9	Validity check - frequency of a variable being selected	54
3.10	Validity check - distribution of R^2 for prediction models	55
3.11	3-hours-blocks breakdown, Left: $\beta_{1,ML}$, Right: $\beta_{2,ML}$	57
3.12	Robustness check - using elasticnet to predict in the first step	59
4.1	Goal-setting households in the experiment	66
4.2	Kernel density and means for violation in goal	67
4.3	The effect of information acquisition costs	84
4.4	The effect of a reminder	85
4.5	The effect of loss aversion	85
4.6	Policy counterfactuals	86
A.1	Survey Email	93
A.2	Survey questionnaire	94
B.1	App usage overtime	103
B.2	Spatial distribution of households in the sample and weather stations . . .	104
C.1	Sampling distribution for Monte Carlo study	109
C.2	Illustration of the gain-loss utility	110

List of Tables

2.1	Sample selection: difference between survey respondents and survey non-respondents	12
2.2	Summary statistics of the sample	13
2.3	Estimation results for Equation (2.1)	16
2.4	Hypotheses for the balance tests	17
2.5	Balance test results	19
2.6	What would happen to prediction error if consumers are attentive for the whole billing cycle	22
2.7	Recall error from previous bill	24
2.8	Estimation results for Equation (2.6) for manual payment consumers	26
3.1	Descriptive statistics for households, pre-treatment energy usage	39
3.2	Descriptive statistics for households, information from their accounts . . .	40
3.3	Descriptive statistics for households, census block information	41
3.4	Descriptive statistics for households, App activity log information	42
3.5	Who sets a goal?	43
3.6	Estimation result for Equation (3.2), daily aggregate usage	46
3.7	“Pre-treatment” prediction error ($kWh_{it} - \widehat{kWh}_{it}$)	52
3.8	Estimation result for Equation (3.4), daily aggregate usage	56
3.9	Estimation result for Equation (3.4), daily aggregate usage, using elasticnet in the first step to predict counterfactual energy usage	58
3.10	Back of the envelop calculation for the mis-statement in the revenue generated by emission abatement	60
4.1	Within-bill dynamics - energy usage	68
4.2	Within-bill dynamics - information acquisition	69
4.3	The effect of an Email reminder on information acquisition	70
4.4	Description of the structural parameters	80
4.5	Summary statistics for structural estimation sample	82
4.6	Estimated structural parameters	83
A.1	Pilot survey results	92
A.2	Balance test results - Wave 1	95
A.3	Balance test results - Wave 2	96
A.4	Balance test results - Wave 3	97
A.5	Robustness checks for for Equation (2.1)	99
B.1	Regression results for diff-in-diff models by 3 hours block	104
B.2	Regression results for ML diff-in-diff models by 3 hours block	105

C.1	Variables description	106
C.2	Simulation baseline parameters	108
C.3	Simulation settings	108
C.4	Comparing Monte Carlo results with actual parameters	109

Chapter 1

Introduction

According to [World Resource Institute \(2018\)](#), residential buildings contributed to nearly 11% of the total 49.4 gigatons of CO₂ emission in 2016. Out of all activities, this is the second-largest contributor to greenhouse gas emissions. Hence, policymakers target the residential energy sector to reduce emissions. To implement effective policies, we must first understand the consumers, for example, their preferences and what incentives they react to.

In standard economic models, individuals choose consumption optimally by maximizing their individual utility, subject to their budget constraints. However, behavioural economics suggests we should also focus on “non-standard” preferences and beliefs, as it has been shown that individuals’ judgements are subject to cognitive and behavioural biases that deviate from full rationality. This thesis explores consumers’ behaviour in the residential electricity industry, and I provide evidence that consumers exhibit inattention and behaviours that are consistent with non-standard preferences.

Intermittent billing, where the consumer is billed periodically, is common across goods and services, including the residential electricity industry. In the absence of information feedback between bills, consumers have to be attentive to their consumption to make fully informed consumption decisions. In the first essay, I study whether consumers are attentive to their energy usage in the absence of any real-time feedback. I ask consumers to predict their electricity bills in a survey. Since consumers pay for the cumulative sum of their usage within the billing cycle, they have a different amount of information (their incurred usage) to predict their bills depending on when they are in their billing cycles. I exploit the random variation in consumers’ billing cycle starting date, so consumers have an exogenous variation in the amount of information to predict their bills, depending on when they respond to the survey in the billing cycle. I expect consumers who respond

to the survey at the end of the billing cycle to have no prediction errors if they are perfectly attentive. I find that prediction errors fall for the first 12 days of the billing cycle, suggesting consumers are attentive to their current billing cycle's usage for 12 days. This is consistent with the previous findings, that receiving a previous bill in the early stage of the billing cycle makes electricity expenditure salient. However, this saliency does not endure for the whole billing cycle.

Electricity is a good that yields immediate benefits and delayed costs due to intermittent billing. In the absence of a commitment device, a present-biased consumer would overconsume relative to what he prefers from an ex-ante perspective. Hence, when the sophisticated present-biased consumer has an option to set a goal as a commitment device, he will use it to curb his overconsumption. In the second essay, I estimate the effect of goal-setting in conjunction with real-time information feedback on energy usage. However, since most goal-setting interventions are voluntary, the common empirical challenge is to address selection into goal-setting. I partner with a local electricity retailer and implement a field experiment that randomly assigns consumers with an App that provides them with real-time feedback. In addition to the real-time feedback, I provide an option for consumers to set an energy usage goal. To address selection, I leverage the richness of my data and use a recently developed machine learning causal inference method that semi-parametrically controls for household-time varying confounders. I find goal-setting generates a \$12 annual saving for a household on average. Without carefully controlling for household-time varying confounders would overstate the effect of goal-setting with real-time information by 64%.

My third essay aims to understand the behavioural mechanism that drives the estimated effect from the second essay. Results from the second essay suggest that goal-setters are reference-dependent. The goal serves as a reference point, so goal-setters' energy consumption decisions are affected by its relative comparison. However, since consumers may not acquire information by using the App, they can be misinformed about their progress towards goal achievement. To understand how information frictions can affect the effectiveness of goals, I propose and estimate a model of decision making. I model a reference-dependent decision-maker that chooses daily energy usage and costly information acquisition, subject to a goal that he sets. The forward-looking decision-maker chooses consumption that maximizes his billing cycle aggregate utility. However, his perceived consumption may differ from actuality in the absence of immediate feedback. He may overcome this discrepancy with costly information acquisition. At the terminal period, he may experience disutility if he consumed more than his goal. Using the data from the field experiment, I estimate this dynamic structural model to recover four behavioral parameters: (i) the disutility of violating a goal (ii) the cost of acquiring information and (iii) biasedness of households' perceived consumption compared to the

actuality and (iv) the effect of a reminder. This allows me to isolate each feature's effect on energy usage. I find that consumers on average reduce their energy usage by 13% to avoid the disutility of missing their goals. Furthermore, while information acquisition cost plays a big role in reducing consumers' interaction with the App, not being fully informed does not affect energy usage. This is because consumers' perceived energy consumption is close to their actual consumption without acquiring real-time information. I speculate this is due to sample selection, as my sample consists of energy conscious consumers who are knowledgeable on energy consumption quantity.

The remainder of this thesis is structured as follows. Chapter 2 to 4 forms the core of this thesis, Chapter 5 concludes. Appendix A to C provides supplementary materials to the corresponding chapters.

Chapter 2

Forgetting and Forecasting Energy Bills

2.1 Introduction

The neoclassical decision-maker considers all relevant information when he makes his decisions. However, the seminal work by [Simon \(1971\)](#) suggests that decision-makers have limited information-processing capacity, so they can be inattentive to information that is relevant to their decisions. This may lead to inefficient outcomes if the error-prone decision deviates from the fully informed one.

Inattention is particularly relevant in the residential electricity industry, where the benefits of electricity usage are experienced today, and the bill comes at the end of the cycle. Hence, in the absence of immediate feedback, consumers have to attentively track their usage over the billing cycle if they want to make fully informed consumption decisions. Not being attentive can lead to bill shocks, where the realized bill at the end of the billing cycle exceeds the consumer's expectation. This is a central policy issue in many markets such as the telecommunication industry in the United States ([CITA, 2011](#)). As addressed by the former President Barack Obama:

"Far too many Americans know what it's like to open up their cell-phone bill and be shocked by hundreds or even thousands of dollars in unexpected fees and charges. But we can put an end to that with a simple step: an alert warning consumers that they're about to hit their limit before fees and charges add up."

Policymakers aim to minimize bill shocks as they may lead to households violating their budget constraints. In these circumstances, households will have to sacrifice their other

consumptions or savings to pay for the unexpected component of their bills. Low-income households, who are relatively less liquid, may even have to take on loans which can lead to financial stress and many other social issues.

With the constant improvement in information technology, almost half of all U.S. electricity customer accounts now have smart meters, a technology that tracks consumers' electricity consumption in real-time. By the end of 2016, U.S. electric utilities had installed about 71 million advanced metering infrastructure smart meters, covering 47% of the 150 million electricity customers in the United States (Mey and Hoff, 2017). While firms have consumers' high-frequency electricity consumption information, investing in the technology to deliver this information to consumers can be costly. Whether an investment can be beneficial depends on the extent to which consumers are attentive to their consumption within the billing cycle in the absence of notifications and feedbacks. If consumers are attentive to their energy consumption without feedback and notification, investing in the technology to deliver this information can be socially wasteful, as firms are delivering redundant information that consumers already know. Moreover, even if consumers are inattentive to their consumption, the optimal timing of notifications will depend on the dynamics of inattention within the billing cycle.

This paper aims to test whether consumers are attentive to their electricity consumption in the residential electricity industry and examine whether attention varies over the billing cycle. I survey consumers to predict their forthcoming electricity bills. Together with consumers' actual bills, I construct a consumer-level prediction error. Since consumers pay for their cumulative electricity usage within the billing cycle at the end of the cycle, more information on what they are going to pay (i.e. their already incurred usage) is revealed progressively within the billing cycle. I exploit the random variation in consumers' billing cycle starting date with respect to the survey email date. Consumers are exogenously exposed to different amounts of useful information that could improve their predictions, depending on when they respond to my survey in the billing cycle. I assess whether the exposure to more information will improve their predictions, as consumers who answer the survey at the end of the billing cycle could provide an error-free prediction of their forthcoming bill if they were attentive to their electricity usage throughout the billing cycle.

I document that absolute prediction errors of forthcoming bills decrease for the first 12 days of the cycle. This is consistent with previous findings in the literature, that consumers exhibit attention cycles within the billing cycle. Furthermore, I find that prediction errors rise after the 12th day, and stabilize at a constant level of \$40. I find that if the attention from the first 12 days extends throughout the billing cycle, prediction errors would reach zero at the end of the billing cycle. This implies that the attention cycle

gives rise to bill shocks. Lastly, contrary to previous studies, I find that consumers who use automatic payment initially have a lower prediction error than consumers who pay their bills manually. This implies automatic bill payment does not reduce the salience of electricity consumption.

2.1.1 Related literature

Previous literature has shown that consumers are inattentive to information in various contexts such as sales tax ([Chetty et al., 2009](#)) and product attributes ([Sallee, 2014](#)). Within the residential electricity industry, consumers are inattentive to complex pricing ([Ito, 2014](#)), the exact amount that they pay ([Keefer and Rustamov, 2015](#)), alternative retailers ([Hortaçsu et al., 2017](#)) and pricing schedule ([Fowlie et al., 2017](#)). Studies show consumers in the U.S. spend on average 7 minutes per year to interact with their electricity provider ([Guthridge et al., 2012](#)), indicating they allocate little of their time and attention to their utility matters.

The closest studies to this paper are [Sexton \(2015\)](#) and [Gilbert and Zivin \(2014\)](#). In [Sexton \(2015\)](#), the author shows that consumers increase their electricity consumption after they enrol in automatic bill payment, as they forgo the inspection of their recurring bills thus reducing the price salience when they use electricity. In regard to the dynamics of attention, [Gilbert and Zivin \(2014\)](#) show that consumers exhibit a consumption cycle within the billing cycle as they decrease their electricity consumption after they receive their bills. [Allcott and Rogers \(2014\)](#) find similar results, as consumers exhibit ‘action and backsliding’ after they receive a home energy report that utilizes a peer comparison. They argue such consumption cycles arise as a result of dynamic salience - electricity expenditures are only salient to consumers when they receive a bill in the early stage of the billing cycle, but they return to over-consumption as their attention to the usage fades overtime.

Unlike previous studies, where they indirectly infer attention (or lack thereof) through surges in consumption, I use a survey to elicit consumers’ prediction as an alternative measure of attention, providing additional evidence to this salience cycle. Furthermore, since my test is based on whether consumers incorporate their already incurred usage into their predictions, I can assess the implication of the salience cycle on bill shock that previous studies are unable to.

The remainder of the paper is structured as followed: Section [2.2](#) provides a thought experiment on how I assess consumers’ attention. Section [2.3](#) describes the data that I analyse and how I exploit the exogenous variation in consumers’ billing cycle starting date to investigate consumers’ forecasting accuracy. Section [2.4](#) describes the empirical

models that I estimate and discusses the estimation results. Lastly, I conclude in Section 2.5.

2.2 A thought experiment

Suppose we ask two similar consumers to forecast their forthcoming bills at different times in their billing cycles, where the bills are the cumulative sum of the incurred electricity usage within the billing cycle¹. These consumers do not have any method to check their usage. We ask the first consumer at the start of his billing cycle and the second consumer at the end of his billing cycle.

The first consumer who forecasts at the start of their billing cycle has minimal usage information on the current billing cycle. He can rely on information such as his last bill and the weather forecast to predict how much he will have to pay at the end of the cycle. However, he is likely to make a forecast error unless he can perfectly forecast his usage for the remainder of his billing cycle.

If the second consumer is perfectly attentive to his daily usage, then I expect him to have a perfect, error-free “forecast”. This consumer can cumulate his overall usage and report the sum as his “forecast”, as he will be billed based on how much he used within the billing cycle. In contrast, if the second consumer is not perfectly attentive to their usage that has already been incurred, then their forecast will not be error-free.

Intuitively, consumers at different times in their billing cycles are exposed to different amounts of available information to forecast their forthcoming bills. Consumers at the start of the billing cycle can rely on external information (e.g. weather forecasts) to forecast their forthcoming bills. As time progresses in the billing cycle, while consumers still possess the aforementioned information, this information becomes less relevant as the current cycle’s incurred usage becomes the most relevant information. In other words, consumers have a weakly larger information set to forecast their incoming bills later in the billing cycle. Since they have no tool to check for the exact amount of usage that has been incurred, whether they have a strictly larger information sets will depend on whether they were attentive to their usage throughout the billing cycle. If consumers are perfectly attentive, then consumers have a strictly larger information set to forecast towards the end of the billing cycle. Hence, I expect the first consumer’s forecast error to be higher than the second consumer’s. In contrast, if consumers are not attentive, then their information sets do not expand as time progresses. Therefore, I can infer attention to

¹The billing cycle does not necessarily have to align with the calendar month.

electricity consumption by comparing the forecast errors between consumers at different times in the billing cycle.

As [Gilbert and Zivin \(2014\)](#) suggests, the relationship between the day of the billing cycle and forecast error does not have to be linear due to the dynamics of inattention. If the current billing cycle's electricity expenditures are only salient to the consumer after they receive their previous bills, then the consumer's information set will expand for the period where electricity expenditures are salient. This causes forecast errors to reduce at the start of the billing cycle when electricity expenditures are salient, as consumers can incorporate their incurred usage into their forecast. However, this salience effect may not endure for the whole billing cycle. Hence, I expect forecast errors to be non-decreasing after the salience effect from receiving the previous bill is no longer present.

2.3 Data

2.3.1 Background and context

The research context for this study is the residential electricity industry in Victoria, Australia. In 2006, the Victorian government decided to mandate a compulsory roll-out of smart meters to replace the old-style "accumulation meters" that simply register households' total electricity use ([Australian Energy Market Commission, 2007](#)). By 2016, they completed the rollout, and all households and small businesses in Victoria upgraded their meters to a smart meter ([Australian Energy Regulator, 2016](#)). Despite having this infrastructure, electricity distributors and retailers are not obligated to provide this information to their consumers. Nevertheless, many retailers and distributors voluntarily provide their consumers with their smart meter information ([AGL Energy, 2020](#); [Origin Energy, 2020](#)). As for the retailer that I partnered with, the company was experimenting with this new technology, so my sampled consumers do not have access to their smart-meter data.

Residential electricity is a suitable context to study inattention for the following reasons. First, electricity bills represent a large proportion of households' disposable income. The average Australian household spends \$39 per week on electricity and gas for their homes ([Australian Bureau of Statistics, 2013](#)). That is equivalent to 10% for low-income households and 3% for high-income households. Second, electricity bills are volatile, and the difference can be up to 300 kWh² in different seasons (see [Australian Energy Regulator, 2017](#), Figure 5.8), so consumers can potentially experience large bill shocks. Hence, consumers have an incentive to pay attention to their electricity consumption to avoid a bill

²This is approximately \$20 dollars per month in a back of the envelope calculation.

shock as it is a large proportion of their household expenditures. However, bill shock is an issue in the residential electricity market in Australia. Bainbridge (2014) finds 1 in 8 Australians cannot afford to pay their electricity bills. One of the contributing factors is that individuals receive the bill some time after usage finishes (say, several days after a monthly billing period completes), they may lose a sense of what they did during that month to suddenly run up such a huge bill (Werdiger, 2012). Hence, avoiding bill shocks is one of the incentives for consumers to be attentive to their electricity usage throughout the billing cycle.

2.3.2 Survey design and data source

My analysis utilises billing data, survey data, and consumer demographic data at the census block level. I partnered with a local electricity retailer and surveyed 70,000 consumers in three waves in November, 2017, June 2018 and July 2018. The energy retailing industry in Victoria is dominated by three main companies, accounting for more than 60% of the customer. The partnering company is a tier below the “big three”.

I sent consumers an Email titled “*Guess your electricity use to go in the draw*”. Upon receiving the Email, consumers can click on a link that directs them to complete the survey. I chose the invitation sample such that the sample characteristics such as payment method, credit rating, and average bill are representative of all customers by the retailer. The survey response rate was 14%. For consumers who responded, 80% of them responded within 3 days.

I asked consumers to provide some demographic information and to forecast their forthcoming monthly bill. The question specific to bill forecasting is:


To the nearest dollar, what do you expect your electricity bill to be next month?

Together with the actual bill of the consumer at the end of the billing cycle, I construct a consumer-level forecast error by taking the difference between the forecast and the realized amount. The retailer calculates the bill amount based on the consumers’ electricity usage and a fixed daily usage cost, subtracting concessions, early payment discounts, and solar feed-in. Bills are typically distributed to consumers 2-3 days via mail or Email (the consumer’s option) after the billing cycle ends. Consumers can use automatic bill payment or pay their bills manually.

I attach an example of an electricity bill in Figure 2.1. The bill includes a breakdown on how the billed amount is calculated, and the duration of the bill. Hence, the consumer is informed on what is being billed and when it covers. It is worthwhile to note that

consumers in the sample do not have any way to check their electricity consumption within the billing cycle³.

FIGURE 2.1: A sample of an electricity bill



I provide incentives to consumers to respond to the survey. However, to avoid consumers changing their energy consumption after the survey to match their forecast, I do not incentivise for prediction accuracy. To assess whether this affects consumers’ predictions, I conducted a pilot survey with another sample of consumers prior to this survey to ensure consumers are just as accurate without any incentives for prediction accuracy. I implement a binary lottery procedure treatment with quadratic scoring rule (Harrison et al., 2014) with the pilot survey, and the results show that there are no statistical differences in consumers’ forecast accuracy with and without the incentive. I include the results for the pilot survey in Appendix A.1.

³Some distributors provide portals for consumers to check for their electricity consumption within the billing cycle (AGL Energy, 2020; Origin Energy, 2020). I asked households whether they have access in the survey, and I exclude all households who have access to these technologies.

Consumers are eligible to enter a draw upon completing the survey, where they can win a prize that is valued at approximately \$300 Australian dollars, and I select three winners from the draw.

In addition, I asked consumers to provide their household characteristics such as number of household members, type of house, number of bedrooms, and more⁴. I attach the complete survey in the Appendix A.2. Furthermore, I obtain consumers' payment method, credit rating and bill payment history from the retailer. This is crucial as the bill contains information on the date the bill starts, ends, issued, and paid. I use this to construct the number of days since the billing cycle started when the consumer responds to the survey.

To control for consumer observable characteristics, I collect census block-level demographic information such as the proportion of internet users within the region and socioeconomic index (SEIFA) from the Australia Bureau of Statistics and match it to consumers in my sample. The lowest level of aggregation available is at the "Statistical Area Level 1" (SA1) census tract, which typically contains 250 homes.

2.3.3 Sample selection

Survey sample selection can potentially bias my results, as consumers who respond to my survey are more likely to be engaged with their utility matters, and hence, more likely to be attentive to their electricity consumption. As a result, I may conclude that consumers are attentive to their electricity usage, but this may be driven by my biased sample that only consists of attentive consumers. I assess this by comparing observable characteristics between survey respondents and survey non-respondents.

Table 2.1 presents the difference in means test for the survey respondents and survey non-respondents. I find respondents and non-respondents to be statistically different. Surprisingly, I find that survey respondents are less likely to pay their bills manually as opposed to using automatic payment. Since consumers forgo the inspection of their recurring bills as they enrol into automatic payment, contrary to my ex-ante belief, survey respondents are less engaged with their electricity payment compared to survey non-respondents. I also find that while survey respondents have an average \$10 lower historical bill before the survey than non-respondents, this may indicate survey respondents to be more conscious and attentive with their energy usage.

⁴It was an error in the send-out of the survey. I intended to send surveys without the household characteristics questions to consumers who already gave us this information previously. However, this was sent to the wrong consumers who had not provided this information yet.

TABLE 2.1: Sample selection: difference between survey respondents and survey non-respondents

	(1)	(2)	(3)
	Respondents	Non-respondents	Differences
<i>From account details</i>			
Credit rating	54.531 (10.606)	59.318 (15.242)	4.787*** (0.161)
Historical mean of bill (before survey)	109.544 (75.793)	119.459 (146.525)	9.915*** (1.637)
Dual fuel customers	0.779 (0.415)	0.699 (0.459)	-0.080*** (0.005)
Pays manually	0.713 (0.452)	0.784 (0.411)	0.071*** (0.005)
<i>From Census</i>			
Expected age	36.927 (5.884)	36.818 (5.855)	-0.108 (0.065)
Expected no. bedroom	3.050 (0.480)	3.013 (0.500)	-0.037*** (0.005)
Proportion of Australian citizen	0.842 (0.096)	0.833 (0.101)	-0.009*** (0.001)
Proportion of houses	0.870 (0.210)	0.851 (0.228)	-0.019*** (0.002)
Proportion of household that lives > national median income	0.187 (0.094)	0.190 (0.100)	0.003** (0.001)
Proportion of household who have internet access	0.806 (0.100)	0.804 (0.102)	-0.002 (0.001)
Proportion of household pays > national median mortgage	0.469 (0.235)	0.481 (0.235)	0.012*** (0.003)
Expected household size	2.586 (0.429)	2.574 (0.439)	-0.011* (0.005)
Proportion of household pays > national median rent	0.589 (0.327)	0.598 (0.320)	0.009** (0.004)
Proportion of renters	0.282 (0.162)	0.297 (0.170)	0.015*** (0.002)
Observations	9805	60154	69959

2.3.4 Summary statistics


I restrict my analysis to monthly billed consumers as I do not have sufficient survey respondents for each day for quarterly billed consumers. I also exclude prediction error (in absolute terms) that are over the 99th percentile to avoid outliers, leaving us with 9,259 observations in my final sample. Table 2.2 presents the summary statistics of the sample. On average, a consumer in the sample lives in a house with three bedrooms and has three people in the household. The average bill is \$110, which is close to the Victorian average of \$1250/year, \$104/month ([The Australian Energy Market Commission, 2017](#)).

TABLE 2.2: Summary statistics of the sample

	No. observation	Mean/Prop	Std dev
<i>(a) Survey data</i>			
Electricity and gas heating	9308	0.247	0.432
Electricity heating	9308	0.268	0.443
Gas heating	9308	0.485	0.500
No of air conditioners	9629	1.231	0.887
Uses central air conditioner	9151	0.185	0.388
Apartment	9330	0.121	0.327
House	9330	0.749	0.434
Terrace	9330	0.012	0.107
Townhouse	9330	0.118	0.323
Owens the property	9361	0.669	0.471
Rents the property	9361	0.331	0.471
Has a pool or a spa	9363	0.102	0.302
Solar hot water	9359	0.190	0.392
Gas water heating	9359	0.737	0.440
Electricity water heating	9359	0.073	0.261
No of bedrooms	9361	3.121	0.910
Household size	9350	2.798	1.327
Experienced bill shock	7852	0.296	0.457
<i>(b) Account data</i>			
Credit rating (out of 100)	9542	54.524	10.433
Total bill (in dollars)	9629	112.132	74.237
Manual payment	9542	0.731	0.443
<i>(d) Census data</i>			
Age	9422	36.960	5.925
No of bedrooms	9411	3.053	0.473
Prop of Australian citizens	9411	0.843	0.094
Prop of houses	9411	0.872	0.206
Median income	9411	0.185	0.094
Prop of internet users	9411	0.804	0.101
Median mortgage payment	9379	0.466	0.235
Household size	9411	2.584	0.429
Median rent payment	9065	0.586	0.328
Prop of renters	9411	0.281	0.160

In order to compare households across different point in time within the billing cycle, the analysis requires variation in time in the billing cycle when the household responds to the survey. Figure 2.2 presents evidence that I do observe consumer across different time in the billing cycle when they respond to the survey.

FIGURE 2.2: Variation in the day of the billing cycle when the consumer responds to our survey.




2.4 Are consumer attentive to their electricity usage?

2.4.1 Graphical evidence

Before presenting the econometrics specification, it is helpful to highlight the graphical relationship between the two key variables from my thought experiment. I plot the relationship between the number of days since the billing cycle started when the consumer responds to the survey and the average root mean square error ($\sqrt{(Prediction - ActualBill)^2}$) of the prediction error within the day with the 95% confidence intervals in Figure 2.3. The figure shows the root mean square error (RMSE) falls for the first 12 days, then rises and stays flat for the rest of the billing cycle, which is consistent with the dynamic salience prediction suggested by [Gilbert and Zivin \(2014\)](#).

FIGURE 2.3: Average RMSE prediction within days



Notes: This figure shows the average root mean square error for each day of the billing cycle.

2.4.2 Econometrics specification

I now estimate the relationship between the number of days since household started their billing cycle when he responds to the survey, and their root mean square prediction error (RMSE) in dollars. As motivated by the graphical evidence and my thought experiment, I use different polynomial functions $f(\cdot)$ to approximate this non-linear relationship:

$$RMSE_{is} = \beta_0 + f(Day_{is}, \beta) + \mathbf{Z}_i \gamma + \tau_s + \varepsilon_{is} \quad (2.1)$$

where Day_{is} represents the day of the billing cycle when consumer i responds to the survey s . For example, if the billing cycle starts on 1st November, and the consumer responds to the survey on 3rd November, $Day_{is} = 3$. $\mathbf{Z}_i$ is a vector of consumer characteristics and census block level characteristics. τ_s is the survey fixed effects, accounting for differences in prediction error for different waves and some surveys do not include consumer demographic questions. I use robust standard error and cluster at the census block level. The turning point of (2.1) is a non-linear combination of estimated parameters, I provide the derivation of the turning point in Appendix A.4. I construct the standard errors of the turning point using the Delta method, where I provide the derivation in Appendix A.4.

TABLE 2.3: Estimation results for Equation (2.1)

	(1)	(2)	(3)	(4)	(5)	(6)
	RMSE	RMSE	RMSE	RMSE	RMSE	RMSE
Day	-0.124 (0.084)	-0.138 (0.097)	-1.570*** (0.327)	-1.574*** (0.377)	-3.221*** (0.773)	-3.061*** (0.881)
Day ²			0.049*** (0.010)	0.048*** (0.012)	0.189*** (0.060)	0.176** (0.068)
Day ³					-0.003** (0.001)	-0.003* (0.001)
Constant	44.169*** (1.498)	-32.663*** (8.302)	51.040*** (2.260)	-25.016*** (8.593)	54.790*** (2.833)	-21.815** (8.653)
Observations	9,353	7,262	9,353	7,262	9,353	7,262
R-squared	0.004	0.072	0.007	0.075	0.007	0.075
Controls	N	Y	N	Y	N	Y
Mean of dep. variable	42.35	43.64	42.35	43.64	42.35	43.64
1st Turning point			16.10	16.23	12.21	12.51
Standard error (first)			0.851	0.980	1.260	1.622
2nd Turning point					27.98	28.59
Standard error (second)					3.602	5.012

Robust standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

Table 2.3 presents the estimation result for Equation (2.1). The linear specification from columns (1) and (2) indicates households' forecast errors do not improve over the billing cycle. However, as motivated in previous sections, it is unable to capture the non-linearity and hence it is mis-specified. Column 3 and 4 from Table 2.3 shows a statistically significant U shape relationship between the absolute forecast error, where the turning point is on the 16th day. The preferred specifications are column 5 and 6 using a cubic approximation, as it allows for the most flexibility. The turning point is estimated to be the 12th day of the billing cycle. I provide different specifications and definitions of $RMSE_{is}$ in Appendix A.5. My results are robust to alternative specifications, and this confirms the prediction error falls for the first 12 days then rises and stays flat for the rest of the billing cycle.

The turning point varies from 12-16 days depending on the specification. This indicates that consumers are attentive to their current billing cycle's energy usage and able to incorporate this information into their expectations for 12 days. This is consistent with Gilbert and Zivin (2014), that the arrival of the bill makes electricity consumption salient to consumers, but this attention only spans for a limited time. However, my turning point is different from theirs, where they find the salience invoked by the arrival of the bill lasts

for 10-12 days. This is possibly driven by sample selection, that only attentive consumers respond to the survey, while [Gilbert and Zivin \(2014\)](#) studies a more general population.

2.4.3 Threats to identification strategy

As motivated by the thought experiment, I use the number of days since the billing cycle starts as a proxy of the amount of information that is exposed to the consumer. In order to have a fair comparison between consumers within the billing cycle, I rely on the fact that consumers' billing cycle starts randomly. In other words, the identification assumption is that the billing cycle timing (and hence, the survey response date) is exogenous to forecast error. I use this natural experiment to differ the amount of information about their forthcoming bill for consumers. Consumers who are surveyed later in the billing cycle have more information about the incoming bill as the consumption has already occurred. However, whether this has an effect on the prediction error depends on consumers' attentiveness to their electricity usage throughout the billing cycle.

To test the validity of this identification assumption, I run balance tests to investigate whether there are systemic differences in consumers' observable characteristics in different times of the billing cycle when they responded to the survey. I run two specifications of regression to investigate this:

$$Z_{is} = \beta_0 + \beta_1 \mathbf{1}\{10 < Day_{is} \leq 20\} + \beta_2 \mathbf{1}\{20 < Day_{is} \leq 31\} + \varepsilon_{is} \quad (2.2)$$

$$Z_{is} = \beta_0 + \sum_{j=2}^{31} \beta_j \mathbf{1}\{Day_{is} = j\} + \varepsilon_{is} \quad (2.3)$$

where Z_{is} represent a consumer characteristic and Day_{is} represent the day of the billing cycle when consumer i responds to survey s . I run this regression for all consumer characteristics variable.

TABLE 2.4: Hypotheses for the balance tests

Specification	Null hypothesis	Alternative hypothesis	Relevant test statistics
(2.2)	$\beta_1 = \beta_2 = 0$	any one of $\beta_1, \beta_2 \neq 0$	F-Statistics
(2.3)	$\beta_2 = \beta_3 = \dots = \beta_{31} = 0$	any one of $\beta_2, \beta_3, \dots, \beta_{31} \neq 0$	F-Statistics

Notes: This table summarize the hypotheses for the balance test I run to test my identification strategy.

I run two specifications of the balance test. In (2.2), I partition the billing cycle into three proportions (0-10 days, 11-20 days and 21-31 days) and create a dummy to indicate whether households responded to the survey in that window. In (2.3), I create dummies

for each day for the independent variable. I test for the joint significance of the slope coefficients to investigate whether there is systemic difference in consumers' observable characteristics in different time in consumers' billing cycle. I summarize the hypotheses of the balance tests in Table [2.4](#).

TABLE 2.5: Balance test results

	10-days interval		Individual days	
	F-statistic	p-value	F-statistic	p-value
<i>(a) Survey data</i>				
Electricity and gas heating	1.653	0.192	1.320	0.110
Electricity heating	1.197	0.302	0.661	0.924
Gas heating	0.337	0.714	0.919	0.596
No of air conditioners	0.852	0.427	1.226	0.181
Uses central air conditioner	0.317	0.728	1.360	0.088
Apartment	2.040	0.130	0.849	0.705
House	3.892	0.020	1.069	0.364
Terrace	0.572	0.564	1.353	0.091
Townhouse	2.201	0.111	1.174	0.233
Owns the property	0.232	0.793	1.388	0.074
Rents the property	0.232	0.793	1.388	0.074
Has a pool or a spa	2.623	0.073	1.399	0.069
Solar hot water	0.631	0.532	0.761	0.827
Gas water heating	0.227	0.797	1.272	0.143
Electricity water heating	0.157	0.854	1.035	0.412
No of bedrooms	2.239	0.107	1.050	0.391
Household size	0.265	0.767	1.227	0.180
Experienced bill shock	0.703	0.495	0.958	0.534
<i>(b) Account data</i>				
Credit rating (out of 100)	2.373	0.093	0.953	0.541
Pays bill manually	2.465	0.085	0.873	0.668
<i>(c) Census data</i>				
Age	0.069	0.933	1.234	0.174
No of bedrooms	0.882	0.414	0.793	0.786
Prop of Australian citizens	2.322	0.098	1.161	0.246
Prop of houses	2.238	0.107	1.312	0.115
Median income	1.300	0.273	1.819	0.004
Prop of internet users	0.959	0.383	1.717	0.008
Median mortgage payment	1.283	0.277	1.702	0.009
Household size	0.613	0.542	0.674	0.915
Median rent payment	0.336	0.714	1.148	0.262
Prop of renters	0.279	0.757	1.038	0.408

Notes: This table shows the results of the balance tests. Column (2) - (3) represents the F-Statistics and p-value of the F-statistics for equation (2.2). Column (3) - (4) represents the F-Statistics and p-value of the F-statistics for equation (2.3). Results shows the day of the billing cycle when the consumer responds to the survey cannot predict the consumer's observable characteristics, implying they are comparable over the billing cycle. As results, this provide assurance to my identification strategies that consumers are exogenously exposed different amount of available information within the billing cycle. I provide balance test for each individual survey waves in Appendix A.3. Results shows the day of the billing cycle when the consumer responds to the survey cannot predict the consumer's observable characteristics in each survey wave too.

Table 2.5 present evidence on the quality of the randomization among the consumer for specification (2.2) and (2.3) respectively. The point estimates for most of the variables are


statistically insignificant, indicating consumers are comparable across the billing cycle. I also present the results for the balance test by each survey wave in Appendix A.3. I also find consumers are comparable across the billing cycle in each survey wave.

2.4.4 Implication of the dynamics of attention

As I motivate earlier, bill shock is a central policy issue in many industries, including the residential electricity industry that I am analysing in this study. Bill shocks happen when the consumer's expected bill deviates from the actual bill that he receives at the end of the billing cycle. Grubb (2014) argues that not being able to track usages within the billing cycle and not being able to incorporate this information into expectation gives rise to this phenomenon.

From the thought experiment, I expect forecast errors to decrease progressively throughout the billing cycle if consumers are attentive to their energy usage. I observe forecast errors cease decreasing after the twelfth day as consumers are no longer attentive to their energy usage after that. Forecast errors are on average \$40 at the end of the billing cycle, giving rise to bill shocks before consumers receive their bills. Would bill shocks be eliminated if consumers are attentive for the whole billing cycle? To access this counterfactual, I extrapolate the fall in the prediction error from the first 12 days to the end of the billing cycle. I illustrate this exercise in Figure 2.4, where the solid line represents the rate of reduction in forecast error for the first 12 days of the billing cycle, and the dotted line represents the hypothetical reduction in prediction errors if consumers remain attentive for the remainder of the billing cycle as if they are in the first 12 days. I choose the 12th day to extrapolate because I find that prediction error falls for the first 12 days then rises and stays flat for the rest of the billing cycle in Table 2.3.

FIGURE 2.4: What if consumers are attentive for the whole cycle?



Notes: This figure shows the extrapolation of RMSE if the consumers extend their attention for the rest of the billing cycle. I observe RMSE on average reaches zero in this counterfactual scenario.

I verify the graphical evidence using the following regression

$$RMSE_{is} = \beta_0 + \beta_1 \mathbf{1}[Day_{is} \leq 12] + \beta_2 \mathbf{1}[Day_{is} \leq 12] \times Day_{is} + \beta_3 Day_{is} + \mathbf{Z}_i \gamma + \tau_s + \varepsilon_{is} \quad (2.4)$$

where $\mathbf{1}[Day_{is} \leq 12]$ is a dummy variable that equals one if household i responded to the survey s in their first 12 days of the billing cycle. Other variables are the same as the previous specification. β_2 represents the magnitude of the reduction in prediction errors for the first 12 days of the billing cycle. The aim of this regression is to understand how much do prediction errors reduce when consumers are attentive, so I can extrapolate this to the end of the billing cycle as if the consumer is attentive for the whole month.

TABLE 2.6: What would happen to prediction error if consumers are attentive for the whole billing cycle

	(1)	(2)
	RMSE	RMSE
1[Day $\leq$ 12]	15.691*** (4.592)	14.365*** (5.303)
1[Day $\leq$ 12] $\times$ Day	-2.121*** (0.330)	-2.080*** (0.376)
Day	0.186 (0.178)	0.119 (0.204)
Constant	38.230*** (3.970)	-36.767*** (9.068)
Observations	9,353	7,262
R-squared	0.009	0.077
Controls	N	Y

Robust standard errors in parentheses

*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Notes: This table shows the estimation result for Equation (2.4). I fit a linear trend for the first 12 days of RMSE and extend. I observe RMSE decreases \$2.12 per day for the first 12 days of the billing cycle.

Table 2.6 shows that the reduction in prediction errors is mainly driven by the first 12 days of the billing cycle, where an additional day of information reduces the RMSE of the prediction by \$2.12. If I extend this reduction for the rest of the billing cycle, the prediction error will be \$0⁵ at the end of the billing cycle. In other words, if consumers are fully attentive for the rest of the billing cycle and behave as if they are in the first 12 days, they could perfectly predict their forthcoming bill at the end of the billing cycle.

My result is suggestive that the dynamics of inattention is a potential reason that gives rise to bill shocks. If consumers are attentive for the whole billing cycle, bill shocks may not necessarily arise, as what consumers expect to pay aligns with what they have to pay, despite the lack of feedback within the billing cycle. However, since the salience that is invoked by receiving the previous bill does not endure for the whole billing cycle, there is a misalignment in the expected payment and the actual payment before consumers receive their bills. This is consistent with the modelling assumption made by the theoretical

⁵Calculation: $38.230 + 15.691 - (-2.121 \times 31) = -9.709$. Since RMSE cannot be negative, I take this to be 0.

work in [Grubb \(2014\)](#). However, my results suggest that the author's two-periods model may not be able to capture this non-linearity in attention.


2.4.5 Investigating the use of available information

Previous bills are a readily available source of information for consumers to predict their current bills. That can be useful for consumers in the early parts of the billing cycle who have not incurred a lot of usage for the current bill. Hence, an alternative explanation of the U-shape relationship is that the only information households use to predict their next bill is their past bill, and consumers' ability to recall the bill either decline overtime or the decline is the composition effect of people paying their bills at random points within the 2-week payment periods. To investigate this alternative channel, I asked consumers to recall their last monthly bills in the survey. The question specific to bill recalling is:

To the nearest dollar, what was your last monthly electricity bill?

I plot the RMSE of the recall error in dollars i.e. $RMSE = \sqrt{(Recall_{is} - ActualBill_{is})^2}$, on last cycle's bill on the number of days since their last bill was issued (Day_{is}) in [Figure 2.5](#).

FIGURE 2.5: Average RMSE for recall by different days



Notes: This figure shows the average root mean square error for each day of the billing cycle.

I verify this using the following regression

$$RMSE_{is} = \beta_0 + \beta_1 Day_{is} + \mathbf{Z}_i \gamma + \tau_s + \varepsilon_{is} \quad (2.5)$$

where $\mathbf{Z}_i$ is a vector of consumer characteristics and census-block level characteristics and τ_s is survey fixed effects. I use robust standard error. The difference between (2.1) and (2.5) is that bills from last cycle do not change within the billing cycle as all consumption (from the last billing cycle) has occurred. While forecast error (for the current billing cycle's bill) can fall within the billing cycle due to additional information, recall error should not. Hence, we would not expect a U-shape relationship from (2.5). If consumers actively check their previous bill when they answer the survey, we would expect β_0 and β_1 to be zero, because all consumers can access to their previous bill regardless of how many days since they received their previous bill.

TABLE 2.7: Recall error from previous bill

	(1)	(2)	(3)	(4)
	RMSE	RMSE	RMSE	RMSE
Day	0.005	-0.249	-0.227	-0.140
	(0.531)	(0.489)	(0.663)	(0.645)
Day ²			0.010	-0.005
			(0.028)	(0.021)
Constant	58.481***	-37.071	59.276***	-37.517
	(7.719)	(27.494)	(7.118)	(26.987)
Observations	9,259	7,334	9,259	7,334
R-squared	0.003	0.009	0.003	0.009
Controls	N	Y	N	Y
Mean of dep. variable	58.55	54.83	58.55	54.83
Turning point($-\hat{\beta}_1/2\hat{\beta}_2$)			11.74	-14.94
Standard error			27.19	122.8

Robust standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1


Results from Table 2.7 shows that (i) the intercept, representing consumers who just received their previous bill is non-zero and (ii) the absolute recall error is not a function of days since the bill was issued increases. These suggest consumers do not use their previous bill to make their forecast, and this behaviour does not vary within the billing cycle. This implies consumers are not only inattentive to consumption that they have

already made within the billing cycle, but also readily available information such as their previous bills.

2.4.6 Does automatic payment matter?

Sexton (2015) shows that consumers increase their electricity usage by 4% after they enrol in an automatic bill payment program as a result of the price becoming less salient. I can test this theory using my data, where we expect the prediction error to reduce more and for a longer duration as the billing cycle progresses for consumers who pay manually. This is because consumers who pay automatically may not inspect their bills when they pay. I plot the average RMSE of the prediction separately by different payment methods in Figure 2.6.

FIGURE 2.6: Average RMSE within different days by payment method



Notes: Left panel represents automatic payment, right panel represents manual payment

I verify the graphical evidence by estimating the following equation

$$RMSE_{is} = \beta_0 + f(\text{Day}_{is}, \text{ManualPayment}_i, \beta) + \mathbf{Z}_i\gamma + \tau_s + \varepsilon_{is} \quad (2.6)$$

where ManualPayment_i is a dummy variable that indicates whether consumer i pays his bill manually as opposed to using automatic payment. Other variables have the same definition as (2.1). $f(\cdot)$ is a polynomial function to approximate the non-linear effect of

Day_{is} , and whether this non-linearity differs for the two subgroups⁶. To examine how RMSE changes throughout the billing cycle differently for the two subgroups, I test for the difference in the first turning point between them. I provide the derivation of the test statistics in Appendix A.6.

TABLE 2.8: Estimation results for Equation (2.6) for manual payment consumers

	(1) RMSE	(2) RMSE	(3) RMSE
Day	0.055 (0.162)	-1.936*** (0.515)	-3.091** (1.308)
Day ²		0.066*** (0.016)	0.163 (0.106)
Day ³			-0.002 (0.002)
Day × Manual payment	-0.273 (0.188)	0.607 (0.655)	0.125 (1.607)
Day ² × Manual Payment		-0.028 (0.021)	0.015 (0.128)
Day ³ × Manual Payment			-0.001 (0.003)
Manual payment	9.697*** (3.206)	5.054 (4.648)	5.995 (5.682)
Constant	36.989*** (2.644)	46.845*** (3.721)	49.574*** (4.455)
Observations	9,246	9,246	9,246
R-squared	0.006	0.009	0.009
Controls	N	N	N
Mean of dep. variable	41.99	41.99	43.64
1st Turning point (automatic payment)		14.63	12.60
Standard error (automatic payment)		1.219	2.255
1st Turning point (manual payment)		17.63	9.342
Standard error (manual payment)		1.382	26.99
Test for same turning point p-value		0.103	0.909

Robust standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

I document two facts from Figure 2.6 and Table 2.8. First, the estimate for the manual payment variable is statistically insignificant. This suggests the level of prediction errors do not differ between the two subgroups throughout the whole billing cycle. Second, prediction errors reduce as consumers are exposed to more information, regardless of

⁶For example, for a quadratic specification, I include the following regressors: Day_{is} , Day_{is}^2 , $Day_{is} \times ManualPayment_i$, $Day_{is}^2 \times ManualPayment_i$, $ManualPayment_i$.

payment methods. I find the first turning point of the polynomial approximate to statistically indistinguishable between the two subgroups. My result is not consistent with the [Sexton \(2015\)](#)'s finding, as I find no evidence that the enrolment in automatic bill payment reduces the salience of electricity consumption.

2.5 Conclusion

In this paper, I test whether consumers are attentive to their electricity consumption in the absence of any immediate feedback. I survey consumers and ask them to predict their bills, and I assess their accuracy by comparing their predictions to the actual bill. I exploit the natural experiment that consumers' billing cycle starting date starts randomly, so they have a different amount of useful information to predict their bills.

By comparing consumers across the billing cycle, I find that prediction errors fall for the first 12 days, implying consumers are attentive to their electricity consumption in the absence of real-time feedback. This is consistent with a dynamic salience story that is proposed by [Gilbert and Zivin \(2014\)](#), as the arrival of the bill temporarily raises the consumers' attention on their electricity consumption as it becomes more salient. However, the attention invoked by receiving a bill only lasts for 12 days, and this does not depend on payment methods. My results are also suggestive that the dynamics of inattention give rise to bill shocks. If I extrapolate the reduction in forecast errors from the first 12 days of the billing cycle to the end of the billing cycle, consumers can perfectly predict their forthcoming bills, thereby eliminating bill shocks.

An alternative story that could explain my empirical results is that the prediction error could be due to consumers implicitly setting goals to reduce their expensive bills. I could not rule out this explanation with the current survey design. Hence, I direct future research to investigate this possible channel.

Reminders are useful to curb forgetfulness, as they bring a particular decision or task to recipients' attention. It has been to be effective in various domains ([Calzolari and Nardotto, 2017](#); [Karlan et al., 2016](#), for example) to induce behavioural changes. However, [Damgaard and Gravert \(2018\)](#) find reminders give rise to non-negligible nuisance costs on the consumer. Hence, the timing and target of the information is crucial so the value of the reminder can be maximized while the cost of the nuisance incurred by it is minimized.

Policymakers can use reminders in this context to raise consumers' attentiveness and salience of electricity consumption. However, they should be used with caution due to the nuisance cost that they incur. My empirical results shed light on the optimal timing of reminders in this context. Policymakers should not send reminders to consumers when

consumers are still attentive to their electricity consumption, as the role of reminder is redundant. However, once consumers are no longer attentive to their consumption, a reminder can act as a previous bill and it can make energy usage salient again, and this can potentially alleviate bill shocks.

Chapter 3

Machine learning from households about setting an energy goal

3.1 Introduction

As the impact of climate change becomes more apparent, policymakers implement behavioural interventions in the residential energy sector to encourage energy conservation as a complement to conventional policies (see [Abrahamse et al., 2005](#), for a review). For example, the Victoria state government in Australia mandates energy companies to send consumers sample energy fact sheets that utilize a peer comparison¹. With the rise of smart-metering technologies, firms have access to consumers' real-time energy usage information. This allows them to provide real-time feedback to their consumers, aimed to steer them into better-informed consumption decisions. Lately, many programs offer consumers the option to choose a personal energy conservation goal in conjunction with real-time information feedback due to its ease of implementation and scalability.

Theoretically, it has been shown that goals can serve as a commitment device and a reference point for time-inconsistent individuals to counter their present-biasedness ([Heath et al., 1999](#); [Hsiaw, 2013](#)). While this is effective empirically in various domains, despite the early study by [Becker \(1978\)](#) and a recent study by [Harding and Hsiaw \(2014\)](#), the causal evidence of goal setting on energy conservation is unclear. Understanding the causal effect of self-chosen and non-binding goals on energy usage is important before any large-scale implementation as [Allcott and Kessler \(2019\)](#) show, nudges impose a non-negligible social cost on individuals, particularly in the residential energy industry.

¹<https://www.esc.vic.gov.au/electricity-and-gas/information-consumers/getting-best-energy-offer#toc--sample-energy-fact-sheet>

There is no consensus on this causal effect because the most common implementations of a goal-setting intervention are voluntary participation programs with self-chosen goals. This is easier to implement and more consistent with liberal paternalism. Some notable examples are financial saving goals (ING, 2018), weight-change calories goals (FitNow Inc., 2020) and daily movement exercise goals (Apple, 2018). The challenge associated with goal-setting with non-experimental data is to address the selection into goal-setting: goal-setters and non-goal-setters are different. Researchers often employ randomized control trials (RCTs) and assign goals to individuals to address this issue. While this the gold-standard in identifying a “clean” causal effect of goals, this can be expensive to implement and often unattractive from an implementation perspective for field partners or organizations².

This paper aims to estimate the effect of real-time information feedback and goal-setting in the residential electricity industry. I obtain propriety data from an electricity retailer, where the company implemented a RCT and randomly assign consumers with an App that provides real-time electricity usage feedback. The App also offers an option for the user to set a budget as an energy usage goal. Through the activity logs with the App, I observe when a treated household sets a goal. Together with their smart-meter data, which records their energy usage in 30-minute intervals, I can observe their usage change after they receive the App, and after they set a goal.

Similar to many common goal-setting interventions already in practice, I do not assign goals to consumers. To address selection into goal-setting, I use a machine-learning difference-in-difference strategy, a recently developed methodology by Burlig et al. (2020). I leverage the high-frequency (30-minutes interval, household-specific) feature of our smart-meter data to train household-hour specific prediction models, to predict counterfactuals that semi-parametrically control for these confounders that vary over households and time. While this method does not fully solve the selection issue, it offers a flexible and data-driven way to control for household-time varying confounders that could bias the estimate.

I find that real-time information feedback alone does not affect energy usage and goal-setting on average corresponds to a \$12 annual saving for a household. The reduction in energy usage is concentrated from 9PM to 9AM. Moreover, I find positive selection into goal-setting. Failure to account for the household-time varying confounders will bias the effect of goal-setting upward by 64%. Conducting a back of the envelopment calculation,

²Besides, while RCTs solve the endogeneity issue, it inevitably introduces experimenter demand effects. Experimental participants may choose to alter their behaviour as they constitute this as the “appropriate behaviour” in response to their assigned goals (Zizzo, 2010).

I find that the government will overpay more than \$82,000 annually for a state-wide goal-setting emission abatement project if it relies on estimates that do not account for selection in goal-setting.

Overall, my results are consistent with consumers exhibiting reference-dependent utility, hence energy usage is affected by the goal that they set. In Chapter 4, I study the mechanism behind this estimated effect by proposing and estimating a structural model of energy usage that features reference dependence and costly information acquisition.

3.1.1 Related literature

The behavioral economics literature has shown non-binding goals are effective in various domains such as education ([Clark et al., 2020](#)), the workplace for workers' performance ([Goerg and Kube, 2012](#)) and weight-loss ([Uetake and Yang, 2018](#)). In practice, exogenously imposed goals are rare. Instead, freely-chosen goals with voluntary participation are the more common implementation of goal-setting due to its ease of implementation and scalability. However, with non-experimental conditions, it becomes difficult to attribute the behavioral change to the goal-setting due to the selection into goal-setting: We expect goal-setters and non-goal-setters to be different. Not accounting for the selection can potentially bias the estimate of goal-setting, which may lead to incorrect policy implications. I follow [Burlig et al. \(2020\)](#) to account for time-varying confounders that results in the bias. This method offers a tractable way to select observables and unobservables that correlate with the goal-setting, whereby controlling for them. I believe this method can be widely adopted to various non-experimental goal-setting interventions when high-frequency data with a large number of covariates are available³.

This paper also contributes to the strand of literature that examines the effectiveness of non-price interventions in the residential energy industry. The information programs and goal-setting studies are the most relevant to this paper. Estimates of the energy savings from feedback technologies vary widely, from none to as much as 20% ([Faruqui et al., 2010](#)), depending on the design of the technology and the sample. Real-time information feedback is usually used in conjunction with other interventions. For example, [Jessee and Rapson \(2014\)](#) finds consumers to be more price-sensitive if they are provided with real-time feedback through an in-home display with time-of-use pricing. [Byrne et al. \(2018\)](#) shows real-time feedback with peer-comparison can result in unintended boomerang effects.

³This is not uncommon due to the rise in tracking and information technology. For example, many devices can track individuals' movement in real-time.

As for goal-setting in the residential utility setting, an early framed field experiment by [Becker \(1978\)](#) shows that consumers reduce their electricity usage if they are given difficult goals (20% reduction). The author was not able to give real-time feedback to consumers due to the technology limitations at that time. [Harding and Hsiaw \(2014\)](#) finds that consumers are responsive to self-set, non-binding goals. Consumers who choose realistic goals (between 0% and 15%) reduce their consumption by more than 11% compared to those who set high and unrealistic goals. However, the authors do not claim this effect to be causal due to the selection into their energy saving program. Lastly, [Agarwal et al. \(2017\)](#) studies goal setting with real-time information provision in a field experiment, where they randomly assign shower meters and goals to Singaporean households to conserve water. They find real-time feedback reduces usage by 9-10%, and non-binding goals can achieve a reduction of up to 19-20%. To the best of my knowledge, this is the first paper to provide an estimate of goal-setting in conjunction with real-time information feedback in the residential electricity industry, where I am able to account for possible household-time varying confounders that could bias this estimate.

The remainder of the paper is structured as followed. In section [3.2](#), I introduce and explain the context of the field experiment. Section [3.3](#) provides a conceptual framework on why we would expect the App and goal-setting have an effect on energy usage. Section [3.4](#) describes the data that I use and provides descriptive statistics of the sample. Section [3.5](#) presents results from the using the panel fixed effect approach. Section [3.6](#) presents results from the using a machine learning approach. I provide robustness checks for the machine learning approach in section [3.7](#). In section [3.8](#), I discuss the policy implications of my results. Lastly, I conclude in section [3.9](#).

3.2 The Experiment

3.2.1 Background and context

The research context for this study is the residential electricity industry in Victoria, Australia. The background and context were thoroughly described in section [2.3](#).

The residential electricity industry is a suitable setting to study the effect goal-setting for the following reasons. First, the data provides an accurate measure of usage, information acquisition, and goal-setting. Second, energy usage is adjustable in response to goal achievement. For example, the average reverse cycle air conditioner costs around \$0.60 per hour to run. As a common practice in Australia, households can save up to \$3 if they switch off the air-conditioner and visit the shopping centre or a library for five hours instead ([Laskie, 2019](#)).

3.2.2 The platform

FIGURE 3.1: Sample of the platform

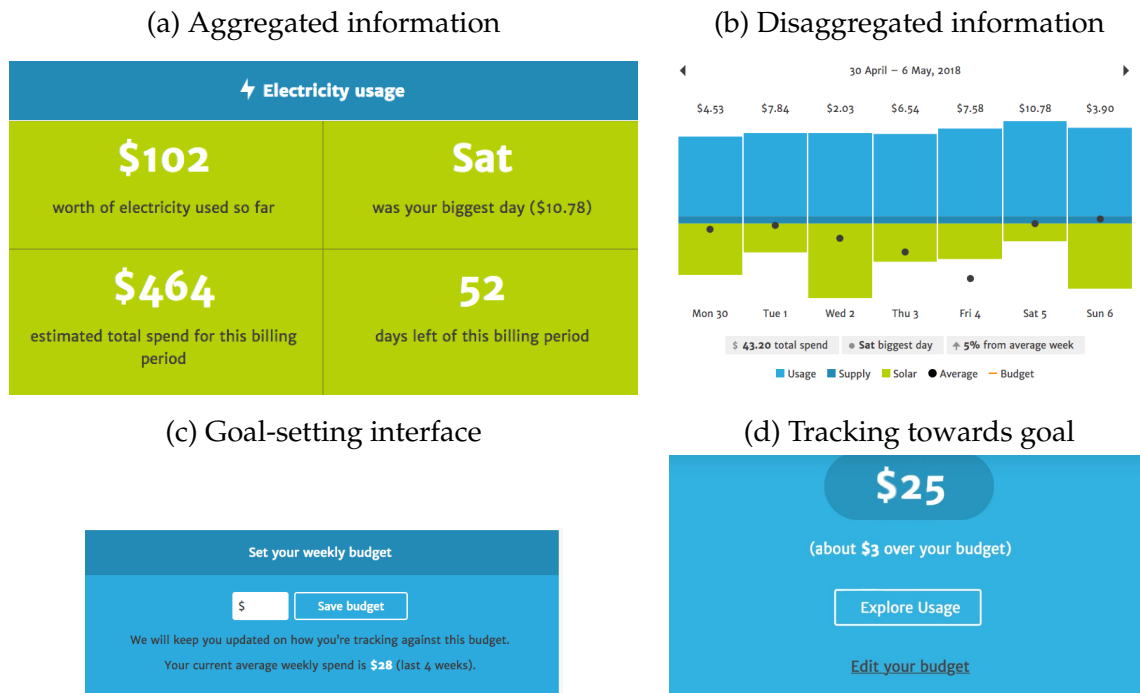


Figure 3.1 is a sample interface of the platform. It provides consumers with their electricity usage information in two ways: a weekly Email and an online portal, where the online portal is available in both the web and a mobile App. The Email and the online portal contain the same information. Hereinafter, I refer the whole platform as “the App”.

From panel (a), I show that the App provides information on billing-cycle aggregated information on (i) past consumption within the billing period (ii) a forecast of usage in the remainder of the billing period (iii) the number of days left in the billing period. This interface is shown in both the Email that the platform sends and the online portal. From panel (b), consumers can view the usage information broken down into different days. From panel (c), the App also has an option for the household to set a budget as goal. This is a self-chosen non-binding goal, meaning I do not provide incentives for household that achieve their goals.

3.2.3 Experimental design

FIGURE 3.2: Graphic representation of the experimental design




Figure 3.2 depicts the sampling procedure. The experiment was conducted by the retailer, under the supervision of Dr. David Byrne. It was rolled out progressively starting in June 2015. The company randomly selected consumers to participate in the experiment base on their baseline characteristics and geographical location to be in an experimental pool. Within the experimental pool, they were randomly assigned to the App progressively, starting in June 2015. Targeted households (hereinafter referred to as “treated households”) were recruited via Email. Upon invitation, the invited household has to activate their App account via a link provided in the Email. A household gains access to the App as they activate their account. The household can opt-in to set a goal as I depict from Figure 3.1 panel (c). If the household does not activate his App account, they *do not* receive any weekly Emails, and hence, they would not be able to set a goal too. Households in the control group were never exposed to App. I collect the data up to February 2019.

Our final dataset consists of 41,775 control households who were never exposed to the App within the experimental period. Within the 29,282 treated households who were invited to use the App, only 802 set a goal within the experiment, accounting for 2% of the treated sample.

3.3 Conceptual framework

The conceptual framework is an application of [Harding and Hsiaw \(2014\)](#). Consumers are present-biased with reference-dependent utility. They encounter a two-stage optimization problem. In stage 0, nature selects whether individuals are given the App, where they gain access to real-time information on energy usage. In stage 1, if they are selected to access the App, they can set an energy usage goal. In stage 2, they choose monthly energy consumption (given the goal if they chose to set it in stage 1).

Consumers do not have access to real-time information feedback in absence of the App, so they make consumption decisions under imperfect information. If consumers can acquire information on their consumption more frequently, they may change their behavior in such a way that aligns more with standard models of consumer demand, as information provision reduces or eliminates biases such as those due to biased beliefs, inattention, and more ([Allcott and Taubinsky, 2015](#)). While I do not have a hypothesis on the direction of this effect, previous studies inform us that estimates of the energy savings from feedback technologies vary widely, from none to as much as 20 percent ([Faruqui et al., 2010](#)).

Hypothesis 1 *The effect of the App without the goal is non-positive.*

In regard to the effect of goal-setting, specifically to endogenously goals, [Harding and Hsiaw \(2014\)](#) have two predictions. First, present-biasedness is necessary to explain why the consumer would set a goal. This is because electricity is a good that yields immediate benefits (e.g., lighting, air-conditioning) and delayed costs due to intermittent billing. In the absence of a commitment device, a present-biased consumer would over-consume relative to what he prefers from an ex-ante perspective when he is not present-biased. Hence, when the sophisticated present-biased consumer with reference-dependent utility has an option to set a goal as a commitment device, he will use it to curb his overconsumption due to his present-biasedness. Hence, following [Harding and Hsiaw \(2014\)](#)'s proposition 1, holding all else constant, without the option to set a goal, I expect households who select into goal-setting to having higher baseline usage than households who do not set a goal since goal-setters cannot use the goal to curb their overconsumption prior to setting one.

Hypothesis 2 *Holding all else constant, goal-setters have a higher baseline usage than non-goal-setters.*

Second, following [Harding and Hsiaw \(2014\)](#)'s proposition 4, if the consumer is not reference-dependent, the selected goal will be unrelated to actual electricity usage. This is because consumers would not experience disutility if they fall short of meeting their

goals if they are not reference-dependent, so they do not have any incentive to change their usage after setting a goal. Conversely, they will consume less electricity if they are reference-dependent as over-consuming relative to their goals will result in disutility. Hence, I expect goal-setting to have a non-positive effect on electricity usage, with the effect strictly negative if goal-setters are reference-dependent.

Hypothesis 3 *If consumers are reference-dependent, goal-setting reduces their energy usage.*

In summary, I have the following predictions. First, the effect of the App that provides real-time information without the goal feature is non-positive. Second, goal-setters exhibit higher baseline electricity usage than non-goal-setters. Third, goal-setting has a non-positive effect on energy usage, and this effect is negative if goal-setters are reference-dependent. In Chapter 4, I propose and estimate a dynamic structural model, where I formalize the effect of goals and information feedback on energy usage.

3.4 Data

The study utilizes data from five sources: smart-meters, App activity logs, consumers' account information database, weather stations, and the census.

3.4.1 Data sources

The outcome of interest is electricity usage. First, I collect consumers' electricity consumption in 30-minute intervals for both the treated and control households from their smart-meters. Second, I collect App usage data for the treatment group. From the real-time activity logs, I observe the timing of each Email sent, clicked, opened, and online portal login⁴ for each household. The activity log also records the goals set by each household (if any). This gives us an accurate measure of the timing of treatments and goal-setting, allowing us to compare usage before and after households are exposed to the App and the setting a goal. I also obtain consumers' account details data for this sample from the retailer. This includes payment methods, credit rating, concession status⁵, billing cycle lengths, whether they bundle their electricity plan with gas, whether they use solar, their payment schedules⁶, location up to a census block, and their actual bills.


⁴The data does not distinguish between online portal login or an App login.

⁵The Victorian government provides a 17.5% discount on electricity usage and service costs to individuals who holds either a pensioner concession card, health care card, or a veterans' affairs gold card. <https://services.dhhs.vic.gov.au/annual-electricity-concession> for more details.

⁶For example, some households are subject to time-of-use payment plans, where the households are subject to a higher per kWh price during the evening.

Electricity consumption is highly affected by weather conditions, so I collected weather data. Since our outcome variable (electricity consumption) is at a 30-minutes frequency, I collect weather data that matches this frequency. I scrape data from Weather Underground, a commercial weather service company owned by IBM that provides real-time weather information on the Internet. They have a weather station network that provides real-time weather information, see Figure 3.3 as an example⁷. They provide variables on temperature, humidity, wind direction, wind speed, wind gust, air pressure, and precipitation. My final dataset consists of 317 weather stations. Since I only have the location of each household up to the census block level, I use the distance-weighted average of three weather stations to proxy for the weather condition experienced by each household at a 30-minute interval. I provide a map that depicts the spatial distribution of households in my sample and weather stations in Appendix B.2.

FIGURE 3.3: An example of a Weather Underground weather station data web-page



Note: I scraped weather data from Weather Underground, a commercial weather service company owned by IBM that provides real-time weather information on the internet. This is a sample of a weather station’s data on a given day from their weather station network.

Lastly, I collect census block-level demographic information such as age and socioeconomic index (SEIFA) from the Australian Bureau of Statistics and match it to households in my sample. The lowest level of aggregation available is at the “Statistical Area Level 1” (SA1) census tract, which typically contains 250 homes.

⁷<https://www.wunderground.com/pws/overview> for details

3.4.2 Who sets a goal?

Table 3.1 presents summary statistics of the pre-treatment electricity usage by treatment groups. Pre-treatment daily energy usage is balanced between the control group and treatment group who do not set a goal since the App was randomly allocated. However, goal-setters have a higher daily energy usage than the treatment non-goal-setters and control group. This is consistent with hypothesis 2, that present biasedness is necessary to explain why the consumer would set a goal. This is suggestive that goal-setters are sophisticated and they are aware of their self-control problem, so they set a goal to curb their overconsumption.

From Table 3.2, I observe statistical significant differences in account characteristics between the control group and treatment group. However, this is likely driven by the large sample size, as they are economically small differences. However, goal-setters are quite different from the other two groups despite the smaller sample size. They are more likely to be on payment concession, monthly billed, and subject to time of use plans. Table 3.3 presents the summary statistics from the census block. Similar to the previous table, I observe statistical differences in census block demographics characteristics between the control group and treatment group, however, it is likely to be driven by the large sample size, as they are economically small differences

Table 3.4 summarizes the App usage for the treatment group for a month. Since the App sends the household weekly Emails from the platform, they usually receive four Emails in a calendar month⁸. The most common way for the household to engage with the App is to open the Email that the App sends, rather than logging into the portal⁹. Households who set a goal are more active than those who do not, even before they set the goal. Households also use the App more after they set a goal.

In an accompanying paper Chan, I analyse the long-run utilization of the App. I use a finite mixture model to identify different patterns of App usage overtime. I identify five subgroups with different utilization patterns overtime: always users, learners, decayers, quarterly users, and never users, with the always users being the largest subgroup in the sample, accounting for 30% of the sample. I attach a summary of the findings in Appendix B.1 for readers who are interested in the long run compliance of the treatment.

⁸This only include Emails from the platform, so it excludes Emails on bills.

⁹The data does not distinguish between online portal login or an App login, so the data is the sum of the two.

TABLE 3.1: Descriptive statistics for households, pre-treatment energy usage

	(1)	(2)	(3)	(4)	(5)	(6)
	Control Group	Treatment w/o goal	Treatment w goal	(1)-(2)	(1)-(3)	(2)-(3)
Daily energy use (kWh)	11.939 (9.397)	11.945 (9.348)	12.328 (9.356)	-0.006 (0.004)	-0.389*** (0.020)	-0.383*** (0.020)
0:00 - 3:00	1.052 (1.308)	1.050 (1.307)	1.098 (1.268)	0.001* (0.001)	-0.046*** (0.003)	-0.047*** (0.003)
3:00 - 6:00	0.898 (1.040)	0.889 (1.003)	0.929 (0.995)	0.009*** (0.000)	-0.031*** (0.002)	-0.039*** (0.002)
6:00 - 9:00	1.360 (1.387)	1.364 (1.369)	1.401 (1.343)	-0.004*** (0.001)	-0.040*** (0.003)	-0.036*** (0.003)
9:00 - 12:00	1.429 (1.581)	1.414 (1.584)	1.464 (1.582)	0.015*** (0.001)	-0.036*** (0.003)	-0.051*** (0.003)
12:00 - 15:00	1.472 (1.708)	1.441 (1.700)	1.513 (1.734)	0.032*** (0.001)	-0.040*** (0.004)	-0.072*** (0.004)
15:00 - 18:00	1.856 (1.988)	1.870 (1.974)	1.952 (2.030)	-0.014*** (0.001)	-0.096*** (0.004)	-0.082*** (0.004)
18:00 - 21:00	2.222 (2.077)	2.250 (2.082)	2.272 (2.074)	-0.028*** (0.001)	-0.050*** (0.004)	-0.023*** (0.004)
21:00 - 24:00	1.650 (1.585)	1.667 (1.591)	1.699 (1.549)	-0.017*** (0.001)	-0.050*** (0.003)	-0.033*** (0.003)
Observations	25813279	8131993	225566	33945272	26038845	8357559

TABLE 3.2: Descriptive statistics for households, information from their accounts

	(1)	(2)	(3)	(4)	(5)	(6)
	Control Group	Treatment w/o goal	Treatment w goal	(1)-(2)	(1)-(3)	(2)-(3)
Credit rating	58.836 (15.107)	59.173 (15.365)	58.728 (14.955)	-0.333* (0.127)	0.108 (0.549)	0.445 (0.560)
On concession (%)	0.210 (0.407)	0.210 (0.407)	0.312 (0.463)	-0.000 (0.003)	-0.102*** (0.015)	-0.102*** (0.015)
Monthly billing cycle (%)	0.087 (0.282)	0.093 (0.291)	0.158 (0.365)	-0.006** (0.002)	-0.071*** (0.010)	-0.065*** (0.010)
Pays bill manually (%)	0.567 (0.495)	0.616 (0.486)	0.652 (0.477)	-0.043*** (0.004)	-0.085*** (0.018)	-0.036* (0.017)
Monthly billing cycle (%)	0.087 (0.282)	0.093 (0.291)	0.158 (0.365)	-0.006** (0.002)	-0.071*** (0.010)	-0.065*** (0.010)
Has solar panel (%)	0.101 (0.302)	0.030 (0.170)	0.077 (0.267)	0.071*** (0.002)	0.024* (0.011)	-0.048*** (0.006)
Dual fuel household (%)	0.660 (0.474)	0.711 (0.453)	0.772 (0.420)	-0.044*** (0.004)	-0.112*** (0.017)	-0.061*** (0.016)
Time of use plans (%)	0.092 (0.289)	0.076 (0.265)	0.100 (0.300)	0.014*** (0.002)	-0.008 (0.010)	-0.024* (0.010)
Non-linear pricing plans (%)	0.563 (0.496)	0.569 (0.495)	0.588 (0.493)	-0.006 (0.004)	-0.025 (0.018)	-0.019 (0.018)
Switched company in 3 years	0.464 (0.499)	0.438 (0.496)	0.318 (0.466)	0.026*** (0.004)	0.146*** (0.018)	0.120*** (0.018)
Observations	41775	28056	802	71057	42577	28858

TABLE 3.3: Descriptive statistics for households, census block information

	(1)	(2)	(3)	(4)	(5)	(6)
	Control Group	Treatment w/o goal	Treatment w goal	(1)-(2)	(1)-(3)	(2)-(3)
Expected age	37.032 (5.711)	37.114 (5.670)	37.055 (6.058)	-0.087* (0.044)	-0.024 (0.205)	0.059 (0.204)
Expected no. bedroom	3.048 (0.474)	3.032 (0.480)	3.036 (0.499)	0.017*** (0.004)	0.012 (0.017)	-0.004 (0.017)
Proportion of Australian citizen	0.843 (0.094)	0.843 (0.095)	0.843 (0.095)	0.000 (0.001)	0.000 (0.003)	0.000 (0.003)
Proportion of houses	0.870 (0.210)	0.864 (0.215)	0.866 (0.221)	0.007*** (0.002)	0.004 (0.008)	-0.002 (0.008)
Average weekly income	741.624 (165.457)	746.908 (170.148)	735.985 (159.465)	-4.863*** (1.281)	5.638 (5.920)	10.922 (6.110)
Household with internet	0.808 (0.101)	0.808 (0.101)	0.804 (0.102)	0.000 (0.001)	0.004 (0.004)	0.004 (0.004)
Average mortgage payment	1959.143 (598.178)	1975.696 (609.445)	1916.369 (578.143)	-14.913** (4.622)	42.774* (21.444)	59.327** (21.934)
Expected household size	2.602 (0.435)	2.590 (0.436)	2.574 (0.441)	0.013*** (0.003)	0.028 (0.016)	0.016 (0.016)
Average rent payment	304.748 (97.960)	306.059 (98.527)	293.708 (96.836)	-0.904 (0.765)	11.040** (3.565)	12.351*** (3.602)
English speaking households	0.754 (0.203)	0.755 (0.201)	0.774 (0.188)	-0.002 (0.002)	-0.020** (0.007)	-0.019** (0.007)
Proportion of year 12 graduates	0.562 (0.155)	0.566 (0.156)	0.549 (0.159)	-0.004*** (0.001)	0.013* (0.006)	0.017** (0.006)
Average education level	10.942 (0.437)	10.950 (0.440)	10.929 (0.420)	-0.008* (0.003)	0.012 (0.016)	0.021 (0.016)
Proportion of employed workers	0.622 (0.105)	0.623 (0.105)	0.625 (0.106)	-0.001 (0.001)	-0.002 (0.004)	-0.001 (0.004)
Observations	41647	27725	795	70588	42442	28520

TABLE 3.4: Descriptive statistics for households, App activity log information

	(1)	(2)	(3)	(4)	(5)
	Treatment Group	Before goal	After goal	(1)-(2)	(2)-(3)
Opens an Email	1.594 (2.320)	2.664 (2.905)	2.902 (3.088)	-1.113*** (0.023)	-0.238*** (0.038)
Click on an Email	0.128 (0.705)	0.514 (1.393)	0.638 (1.800)	-0.402*** (0.007)	-0.124*** (0.021)
Login portal	0.177 (1.027)	0.789 (2.223)	0.971 (2.839)	-0.638*** (0.010)	-0.182*** (0.033)
Any of the above	1.654 (2.444)	2.991 (3.336)	3.262 (3.756)	-1.391*** (0.024)	-0.271*** (0.046)
Email received	3.669 (1.479)	3.833 (1.269)	4.008 (1.129)	-0.174*** (0.015)	-0.174*** (0.015)
Observations	797050	9735	17435	779615	27170

Notes: This Table summarize the App activity over a calendar month. Email received represents the number of Emails the platforms sends to the consumer, it does not include Emails on bills. Opens an Email represents the consumer opening an Email that the platform sends. Clicking on an Email represents the consumer clicking a link on the Email that the platform sends that direct them to the portal. Login portal represents accessing the online portal either through the mobile App or the web.

To formally test hypothesis 2 on whether goal-setters have a higher baseline energy usage than non-goal-setters, I run the following regression

$$kWh_i = \beta_0 + \beta_1 \text{Goal-setter} + \mathbf{Z}_i \gamma + \epsilon_i \quad (3.1)$$

where kWh_i represents the average pre-treatment daily electricity usage for household i . Goal-setter_i is a binary variable that indicates whether household i eventually sets a goal after they are given the App. $\mathbf{Z}_i$ represents a vector of household characteristics. This includes account information such as credit rating, whether the household uses solar and more. I use robust standard errors. From hypothesis 2, holding all else constant, I expect goal-setters to have a higher baseline electricity usage before they are given the option to set a goal. Hence, I expect β_1 to be positive.

TABLE 3.5: Who sets a goal?

	(1)	(2)	(3)	(4)
	Pre-treatment daily usage	Pre-treatment daily usage	Pre-treatment daily usage	Pre-treatment daily usage
Goal-setter	0.849*** (0.254)	0.815*** (0.257)	0.936*** (0.254)	0.873*** (0.268)
Credit rating		0.043*** (0.002)	0.041*** (0.002)	0.045*** (0.002)
Has concession card		-1.399*** (0.060)	-1.384*** (0.059)	-0.960*** (0.067)
Monthly billing cycle		0.393*** (0.089)	0.384*** (0.088)	0.310*** (0.094)
Pays bill manually		-0.477*** (0.064)	-0.532*** (0.064)	-0.297*** (0.069)
Solar user		-0.887*** (0.089)	-0.439*** (0.109)	-0.902*** (0.123)
Dual fuel users		0.606*** (0.074)	-0.073 (0.078)	0.011 (0.098)
Constant	11.543*** (0.026)	9.390*** (0.136)	10.049*** (0.139)	9.526*** (0.156)
Observations	69,900	57,550	57,498	56,102
R-squared	0.000	0.019	0.057	0.310
Household characteristics	N	Y	Y	Y
Contract FE	N	N	Y	Y
Census-block FE	N	N	N	Y

Robust standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

Notes: This presents the results for Equation (3.1), from the most basic to the most stringent specification. The outcome variable is the average daily electricity usage. In column (1), I do not include any household controls. In column (2), I include some observable household characteristics. In column (3), in add contract fixed effect, as households sign different contracts differs in per kWh cost, non-linear pricing, time-of-use pricing and more. In column (4), I add census-block fixed effects, as I observe households' geographic location up to their census-block. I measure all variables at the pre-treatment period.

Table 3.5 presents the results for Equation (3.1), from the most basic to the most stringent specification. The results show that goal-setters have a higher average pre-treatment daily usage than non-goal-setters. This is consistent with hypothesis 2, that present biasedness is necessary to explain why the consumer would set a goal. In the absence of a commitment device (a goal), a present-biased consumer would over-consume relative to what he prefers from an ex-ante perspective when he is not present-biased.

These results highlight the observable differences between goal-setters and non-goal-setters. These differences may be correlated with electricity usage, it is important that I consider selection into goal-setting as a possible threat to econometric identification in this setting. Because I have repeated observations for each household over time, I will employ a panel fixed effects approach, meaning that level differences alone do not constitute threats to identification. For the panel fixed effects results to be biased, there must

be time-varying differences between goal-setters and non-goal-setters that correlate with the timing of goal-setting.

3.5 Panel fixed effects approach

To investigate the effect of the App and goal-setting, I run the following difference-in-difference (DiD) model:

$$kWh_{it} = \alpha_i + \tau_t + \beta_1 \text{App}_{it} + \beta_2 \text{Goal-and-App}_{it} + \epsilon_{it} \quad (3.2)$$


where kWh_{it} is electricity consumption, measured in kilowatt hours, for household i , day t . α_i and τ_t are household and date fixed effect respectively. App_{it} takes on the value 1 if household i is assigned to the treatment group, and if t is after I send the treated household the invitation Email. Goal-and-App_{it} takes on the value 1 after a treated household i sets a goal on the App. I also estimate this relationship for each 3-hour block, highlighting the treatment effect for different time of the day. I use robust standard errors, clustered at the household level to account for arbitrary within-household correlations.

β_1 from Equation (3.2) corresponds to an Intent-to-Treat (ITT) effect of the App alone without a goal on energy consumption¹⁰, as not every treatment household complies. Since households can only set a goal after they receive the App, when Goal-and-App_{it} is one, App_{it} has to be one. β_2 in Equation (3.2) represents the additional effect of the App with the goal, on top of the information that the App provides. From hypothesis 1, I expect β_1 to be non-positive. From hypothesis 3, I also expect $\beta_1 + \beta_2$ to be nonpositive.

I present pre-trends for each 3-hours blocks in Figures 3.4. It plots the differences in the mean (treated - control) in energy usage prior to the experiment, I observe a slight upward pre-trend.

¹⁰Since I do not distinguish the effect between the mobile App or the online portal, we can interpret the β_1 as the intent-to-treat effect of the access to information too.

FIGURE 3.4: Parallel trend assumption - in three hours blocks



Notes: This graph shows the average energy consumption across treatment and control group before and after the treatment. It should be noted that the experiment rollout was progressive, so the red vertical line represents when the first household receives the treatment around 22nd June, 2015.

3.5.1 Results - DiD

I present the estimated coefficient for Equation (3.2) in Table 3.6.

TABLE 3.6: Estimation result for Equation (3.2), daily aggregate usage

	(1)	(2)	(3)	(4)	(5)	(6)
	Daily usage	Daily usage	Daily usage	Daily usage	Daily usage	Daily usage
App	0.565*** (0.051)	0.320*** (0.019)	-0.040* (0.024)	0.565*** (0.051)	0.329*** (0.019)	-0.035 (0.024)
App-and-Goal				0.017 (0.307)	-0.450*** (0.143)	-0.280** (0.142)
Constant	11.949*** (0.030)	12.050*** (0.008)	12.198*** (0.010)	11.949*** (0.030)	12.050*** (0.008)	12.199*** (0.010)
Observations	57,303,515	57,303,515	57,303,496	57,303,515	57,303,515	57,303,496
R-squared	0.001	0.527	0.587	0.001	0.527	0.587
Household FE	No	Yes	Yes	No	Yes	Yes
Date FE	No	No	Yes	No	No	Yes
Mean of dep. variable	12.18	12.18	12.18	12.18	12.18	12.18

Robust standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

Notes: This Table presents the estimation results for equation (3.2). The outcome variable is daily energy usage in kWh. From column (3), I observe the App reduces energy usage by 0.04kWh per day. However, as indicated by column (5), the reduction is driven mainly by goal-setting.

In column (3) of Table 3.6, the App reduces daily energy usage by 0.04kWh, which is equivalent to 0.33% of the pre-treatment mean usage. This represents the overall impact of the App, including the optional goal-setting feature. In column (4) - (6), I further decompose this effect into the access to information only (App) and the access to information with the goal (App-and-Goal). My preferred specification is column (6) with both date fixed effects and household fixed effects. I find the reduction is mainly driven by goal-setting, as the App variable is no longer statistically significant. Goal-setting reduces daily energy usage by 0.28kWh, which is equivalent to 2.3% of the pre-treatment mean usage.


A back of the envelope calculation suggests that the App generates approximately \$4 of saving annually¹¹ but it's not statistically significant. On the contrary, goal-setting is associated with a reduction of consumption of -0.28 kWh daily on top of the real-time information alone, five times the magnitude of the information feedback. A back-of-the-envelope calculation shows that is around \$30 annually.

I also estimate equation (3.2) in 3-hours blocks to highlight heterogeneity of the effects within the day. There are eight 3-hours block in a day: 0:00-3:00, 3:00-6:00, 6:00-9:00,

¹¹A typical household pays \$0.3/kWh. Hence, saving per year is $-0.035 \times 0.3 \times 365 \simeq 4$

9:00-12:00, 12:00-15:00, 15:00-18:00, 18:00-21:00 and 21:00-24:00. I plot the 3-hour blocks regression estimates in Figure 3.5. The estimation results in table forms are attached in Appendix B.3.

FIGURE 3.5: 3-hours-blocks breakdown, Left: β_1 , Right: β_2



Notes: These two graphs show the point estimates and 95% confidence of the estimates for Equation (3.2) by 3-hours block. The left panel represents the effect of the App (β_1), and the right panel represents the effect of the App with goal (β_2). I provide the results in the table format in Appendix B.3.

The App reduces electricity consumption from 12PM to 9PM, peaking at -0.01 kWh/3 hours in the evening. Goal-setting's reduction happens during the evening and midnight from 3PM to 6AM, and it has an average of reduction -0.05 kWh/3 hours during that time, which is equivalent to 3% of the pre-treatment usage mean.

3.6 Machine Learning approach

$\hat{\beta}_2$ in (3.2) may not have a causal interpretation as we suspect there are household-time varying unobservable characteristics that correlate with the decision to set a goal. The ideal identification strategy is to randomly assign the option to set goals to households, however, it was technologically infeasible to implement.

To improve upon the identification strategy from (3.2), we can introduce multi-dimensional fixed effects to control for unobservables that affect energy usage in different dimensions. For example, household-weekday fixed effect controls for unobservable characteristics that affect a specific household's consumption at a particular day of the week. By doing so, we can isolate the effect of goal-setting on energy usage after controlling for these unobservables. However, given the high-frequency of the smart-meter data and the large sample size, there are theoretically an infinite set of model specifications that I can choose

from, making this problem impractical due to its computational burden. To facilitate this, [Burlig et al. \(2020\)](#) offer a new methodology to select the fixed effects. They show that this methodology is more robust than using fixed effects models and matching methods.

This method involves a two-step procedure. In the first step, I predict electricity consumption using the households' pre-App data. For control and non-goal-setting households, I follow [Burlig et al. \(2020\)](#) and randomly assign a pseudo-treatment date between the 20th percentile and 80th percentile of in sample calendar date. The list of predictor candidates for prediction are month dummies, day of the week dummies, holiday dummies, temperature (and squared), humidity (and squared), wind speed (and squared), precipitation (and squared), air pressure (and squared) and the interaction between these variables. In the second stage, I regress Equation (3.2) but with the prediction error (data minus prediction from the first stage) on the left hand side.

Similar to the usual panel fixed effects approach, this machine-learning-difference-in-difference (ML-DiD) method is not immune to all endogeneity problems. My regression estimates are still biased if there exist energy consumption changes that *coincide directly* with setting a goal. For example, if the household undertakes additional energy conserving behavior or upgrade their appliances to higher energy efficiency *at the same time* they set a goal on the App, then my estimates are biased downwards (more negative). Therefore, I can consider my estimates to be an upper bound of the magnitude of the effect of setting a goal on electricity usage.

3.6.1 The first step

In the first step, I predict households' counterfactual electricity consumption. Since there are many predictor candidates, I face a *variable selection* problem, also known as the *bias-variance trade-off*. If I select too few variables (underfit), I will omit good predictors, resulting in poor out-of-sample prediction. However, if I select too many variables (overfit), I may use poor variables that explain noises in the training sample. In this context, I want to select the variables that best predict electricity consumption, so I can sufficiently control what electricity consumption would happen in the post-App and post-goal periods in absence of the App, i.e., predicting for counterfactuals.

The goal of building the counterfactual here is to generate a good prediction. I can estimate different prediction models with trial-and-error. However, this approach is computationally expensive given the size of the dataset and the dimension of the predictor candidates. To address this, machine-learning methods offer a computationally tractable way to predict these counterfactuals in data-driven manner. Similar to [Burlig et al. \(2020\)](#),

I estimate household, hour-of-day specific prediction models as it offers the most flexibility¹². I desire this flexibility as electricity consumption is highly heterogeneous in two dimensions: across households and across time.

Abadie and Kasy (2019) show that LASSO is robust to many data-generating-processes, so I implement LASSO in my baseline specification. In practice, LASSO is a regularized regression that solves:

$$\hat{\delta} = \underset{\delta}{\operatorname{argmin}} \frac{1}{N} \sum_{i=1}^N \frac{1}{2} (y_i - \delta^T x_i)^2 + \lambda \left[(1 - \alpha) \frac{\|\delta\|_2^2}{2} + \alpha \|\delta\|_1 \right] \quad (3.3)$$

The first term is equivalent to OLS's objective function, where y_i is the outcome of interest, that is electricity consumption in kWh for each household's hour in this context. x_i are the extensive list of predictor candidates that I listed earlier. δ is the vector that represents the effects of each variable on the outcome. The second term is the shrinkage component of the objective function; it shrinks coefficients toward zero. LASSO is a special case where $\alpha = 1$ ¹³. $\|\cdot\|_2^2$ and $\|\cdot\|_1$ are the L-2 and L-1 norms respectively. λ is the tuning parameters chosen by the econometrician, a larger λ corresponds to a higher penalty for an additional predictor.

The algorithm splits the pre-App sample into training and testing sets, and it estimates the for $\hat{\delta}$ according to Equation (3.3) for different λ using the training set. Then it chooses λ that minimizes the mean absolute error in the testing set¹⁴. I emphasize that this procedure is repeated for each household hour-of-day, meaning the sets of predictors that are chosen for each household hour-of-day model can potentially be different. I illustrate this process graphically in Figure 3.6 and 3.7.

¹²That is, each household, hour-of-the-day combination can potentially select a different set of covariates and have a different estimate to predict counterfactual electricity usage. This is essentially interacting household fixed effect and hour-of-day fixed with each of the covariate candidates when the algorithm selects the covariates.

¹³ $\alpha = 0$ corresponds to ridge and $\alpha \in (0, 1)$ corresponds to elasticnet. I use these specifications in my robustness checks.

¹⁴I implement this using the `cv.glmnet` package in R.

FIGURE 3.6: Graphical illustration of identification (ideally)





FIGURE 3.7: Graphical illustration of identification (in practice)



Notes: Left figure represents a treated household, right figure represents a control household, where the dotted vertical line represents the pseudo treatment date.

3.6.2 The second step

In step two, I estimate the following equation:

$$kWh_{it} - \widehat{kWh}_{it} = \alpha_i + \tau_t + \beta_{1,ML} App_{it} + \beta_{2,ML} Goal\text{-}and\text{-}App_{it} + \gamma posttrain_{it} + \varepsilon_{it} \quad (3.4)$$

$$\widehat{kWh}_{it} = \hat{\delta}_{it}' x_{it}$$

where $\hat{\delta}_{it}$ is the vector of LASSO estimates from Equation (3.3)¹⁵, and x_{it} is the vector of candidate variables that I describe previously. Consequently, $\widehat{kWh}_{it}$ is the machine-learning produced prediction from step one. $posttrain_{it}$ is a dummy that equals one

¹⁵It should be noted that the LASSO estimates are household-time specific, as I estimate household-time specific models.

during the out-of-sample prediction period. It accounts for the possible bias in the out-of-sample prediction by re-centering the prediction errors around zero. Other variables follows the same definition as Equation (3.2). To remove the effect of outliers due to poor out-of-sample predictions, I only include observations between 1% to 99% percentile of the prediction error ($kWh_{it} - \widehat{kWh}_{it}$). Similar to Equation (3.2), I use robust standard errors, clustered at the household level¹⁶.

The standard difference-in-difference estimator's identifying assumption is that the treated and the untreated households' electricity usage are trending similarly. The key difference for the ML-DiD estimator is that rather than the electricity usage that trends similarly, it is the prediction error that trends similarly between households with and without goals. Intuitively, ML-DiD is "controlling for" the counterfactual consumption - what would have happened to households' consumption in absence of the App.

3.6.3 Validity checks and diagnostics

Before presenting my results, I perform some diagnostics to assess the performance of my prediction from the first step.

If the machine learning algorithm is performing as desired, we would expect the prediction error ($kWh_{it} - \widehat{kWh}_{it}$) to be zero in the pre-App period. Table 3.7 confirms this and the prediction error is statistically indistinguishable from zero, meaning the prediction models performs well in-sample.

¹⁶These standard errors do not account for prediction errors from the first step. However, as [Burlig et al. \(2020\)](#) note, *there is no guidance from the econometrics literature on doing proper inference in this panel machine learning estimator*. They present block-bootstrap standard errors in the paper, and the estimates are very similar to the clustered standard errors.

TABLE 3.7: “Pre-treatment” prediction error ($kWh_{it} - \widehat{kWh}_{it}$)

	(1)	(2)
	Control Group	Treatment Group
<i>3 hours block prediction error</i>		
0:00 - 3:00	-0.000 (0.651)	0.000 (0.667)
3:00 - 6:00	-0.000 (0.529)	0.000 (0.542)
6:00 - 9:00	0.000 (0.700)	0.000 (0.722)
9:00 - 12:00	0.000 (0.913)	0.000 (0.925)
12:00 - 15:00	0.000 (1.076)	0.000 (1.071)
15:00 - 18:00	0.000 (1.246)	0.000 (1.245)
18:00 - 21:00	-0.000 (1.178)	0.000 (1.208)
21:00 - 24:00	-0.000 (0.856)	0.000 (0.887)
Observations	12178951	7967856

Figure 3.8 plots the number of selected covariates for each model against the size of the pre-App sample size. LASSO performs best when the underlying data-generating-process is sparse (Abadie and Kasy, 2019). It penalizes irrelevant variables, meaning that the optimal model for any given household-hour will not include all candidate regressors. Figure 3.8 supports the choice of LASSO, as the number of chosen predictors does not scale linearly with the size of the training set (pre-App period).

FIGURE 3.8: Validity check - number of variables selected




Figure 3.9 plots fraction of model that selects a certain variable by treatment. I observe the selected variables are similar between treated and control households. This adds validity to the ML method, as we expect households to be balanced in the pre-App period. Furthermore, I observe a high level of heterogeneity in the selected variable across models (households hour-of-day). For example., *weekdays* $\times$ *month* dummies are selected by 80% of the models, while *holiday* dummies are only selected less than 10% of the models. Weather variables such as temperature and perception are also frequently selected, which is consistent with my ex-ante expectation that energy usage is highly affected by weather conditions. This highlights the heterogeneity in electricity consumption across households and time which the machine learning algorithm confirms.

FIGURE 3.9: Validity check - frequency of a variable being selected





Figure 3.10's left panel presents a kernel density of sampling distribution of prediction models' R^2 . I further breakdown the distribution into three-hour-blocks in the right panel. Some of the models have very poor R^2 that is close to zero, coming mostly from midnight. This is expected as changes in midnight electricity consumption are usually irregular. Some examples are watching sporting events in different time zones and late night parties, hence harder to predict with the variable list. This will create outliers in prediction errors ($kWh_{it} - \widehat{kWh}_{it}$), hence this justifies my approach to limit our sample to 1th to 99th percentile of prediction error.

FIGURE 3.10: Validity check - distribution of R^2 for prediction models



3.6.4 Results - ML-DiD

I present the estimated results for Equation (3.4) in Table 3.8.

TABLE 3.8: Estimation result for Equation (3.4), daily aggregate usage

	(1)	(2)	(3)	(4)
	Daily usage	Prediction error	Daily usage	Prediction error
App	-0.040*	0.002	-0.035	0.003
	(0.024)	(0.011)	(0.024)	(0.011)
App-and-Goal			-0.280**	-0.115**
			(0.142)	(0.056)
posttrain		0.255***		0.254***
		(0.017)		(0.017)
Observations	57,303,496	49,272,295	57,303,496	49,272,295
R-squared	0.587	0.117	0.587	0.117
$\beta_1 + \beta_2$			-0.316	-0.112
Standard error			0.0142	0.0568


Robust standard errors in parentheses

*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Note: Column (1) & (3) presents the estimation results from the conventional diff-in-diff estimation (same as Table 3.6) for comparison. Column (2) & (4) presents the estimation results from equation (3.4). The outcome variable is the prediction error of electricity usage aggregated at the daily level.

In column (2) of Table 3.8, the App overall increases daily energy usage by 0.002kWh, but it is not statistically significant. In column (4), I further decompose this effect into the real-time feature only (App) and real-time information with the goal (App-and-Goal). I find the reduction is mainly driven by goal-setting. It reduces daily energy usage by 0.115kWh, which is equivalent to 0.95% of the pre-treatment mean usage. A back-of-the-envelope calculation shows this reduction is equivalent to approximately \$12 annually¹⁷. Comparing to the conventional DiD estimation from Table 3.6, the total effect of goal-setting with real-time information ($\beta_1 + \beta_2$) is reduced by almost 64% from -0.316 to -0.112 . In other words, if I ignored selection into goal-setting, I would overstate the causal effect of goal-setting by 64%.

¹⁷The average price/kWh is 0.3. Hence, $-0.115 \times \$0.3 \times 365 = \12.264 .

FIGURE 3.11: 3-hours-blocks breakdown, Left: $\beta_{1,ML}$, Right: $\beta_{2,ML}$ 

Notes: These two graphs shows the point estimates and 95% confidence of the estimates for Equation (3.4) by 3-hours block. The left panel represents the effect of the App (β_1), and the right panel represents the effect of the App with goal (β_2). I provide the results in the table format in Appendix B.4.

I plot the 3-hour blocks regression estimates in Figure 3.5 to understand the how usage changes within the day¹⁸. I observe the conservation is concentrated from 9PM to 9AM, and it has an average reduction of $-0.015\text{kWh}/3$ hours, which is equivalent to 1.2% of the pre-treatment usage mean.

3.7 Robustness checks

My results can be sensitive to the prediction algorithm that I choose in the first step. While the validity checks and diagnostics confirms LASSO performs well, I use an alternative machine learning algorithm to predict counterfactual electricity consumption to examine whether the results are robust to the prediction algorithm choice. In particular, I use an elastic net regularization in step one to predict counterfactuals. Recall the objective function for the regularized regression previously,

$$\hat{\delta} = \underset{\delta}{\operatorname{argmin}} \frac{1}{N} \sum_{i=1}^N \frac{1}{2} (y_i - \delta^T x_i)^2 + \lambda \left[(1 - \alpha) \frac{\|\delta\|_2^2}{2} + \alpha \|\delta\|_1 \right] \quad (3.3)$$

LASSO is a special case where $\alpha = 1$. $\alpha = 0$ corresponds to ridge regression and $\alpha \in (0, 1)$ corresponds to elasticnet. $\|\cdot\|_2^2$ and $\|\cdot\|_1$ are the L-2 and L-1 norms respectively. λ is the tuning parameter chosen by the econometrician, a larger λ corresponds to a higher penalty for an additional predictor. In my robustness check, I choose $\alpha = 0.5$.

¹⁸The estimation results in table forms are attached in Appendix B.4 for interested readers.

I present the estimated coefficient for robustness check in Table 3.9 and the 3-hours-blocked estimates in Figure 3.12. The results are similar to my baseline results, suggesting my results is not sensitive to my chosen prediction algorithm.

TABLE 3.9: Estimation result for Equation (3.4), daily aggregate usage, using elasticnet in the first step to predict counterfactual energy usage


	(1)	(2)
	Prediction error (in kWh)	Prediction error (in kWh)
App	0.003 (0.011)	0.005 (0.011)
App-and-Goal		-0.110* (0.057)
posttrain	0.258*** (0.017)	0.257*** (0.017)
Observations	49,857,220	49,857,220
R-squared	0.119	0.119

Robust standard errors in parentheses

*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

This table presents the estimation results from equation (3.4), using elasticnet as the prediction algorithm in the first step when I predict counterfactual electricity usage. The outcome variable is electricity usage aggregated at the daily level.

FIGURE 3.12: Robustness check - using elasticnet to predict in the first step



Notes: These two graphs shows the point estimates and 95% confidence of my estimates for Equation (3.4) by 3-hours block. The left panel represents the effect of the App (β_1), and the right panel represents the effect of the App with goal (β_2). I provide the results in the table format in Appendix B.5.

3.8 Policy implications

As part of the strategy to reduce CO2 emissions, the Australian government establishes the Emission Reduction Fund (ERF) to provide incentives for organizations and individuals to adopt new practices and technologies that reduce their emissions. Participants can submit projects and take part in an auction that bids for carbon credit units. One carbon credit unit is earned for each tonne of carbon dioxide equivalent (tCO₂-e) stored or avoided by a project, where participants can sell these carbon credit units to generate income. As of December 2020, approximate 1.5% of the total carbon credit units are rewarded to energy efficiency projects ([Clean Energy Regulator, 2020](#)).

As I show earlier, real-time information feedback with goal-setting is effective in reducing households' electricity usage, thereby reducing carbon dioxide emissions. According to the legislation, the abatement calculation should be the difference between a treatment group and a control group (*Carbon Credits (Carbon Farming Initiative-Aggregated Small Energy Users) Methodology Determination 2015. (Cth)*), suggesting it examines a causal reduction. Combining my empirical results with other sources, I perform a back of the envelop calculation on the emission abatement of a state-wide electricity goal-setting program. I examine the discrepancy in the calculation if do not control for household-time varying confounders.

TABLE 3.10: Back of the envelop calculation for the mis-statement in the revenue generated by emission abatement

	DiD	MLDiD	Source
Reduction in electricity usage per day	−0.28 kWh	−0.115 kWh	this study
Carbon dioxide equivalent (tCO ₂ -e)/kWh	0.01 tonnes/10 kWh	0.01 tonnes/10 kWh	^a
Reduction in CO ₂ emission/day	−0.00028 tonnes	−0.000115 tonnes	
No. of residential dwelling in Victoria	2,520,912	2,520,912	^b
Goal-setting participation rate	0.028	0.028	this study
Expected goal-setting participation population	70,059.3050	70,059.3050	
Total reduction in CO ₂ emission/day	−22.139 tonnes	−7.847 tonnes	
Total reduction in CO ₂ emission/year	−8080.640 tonnes	−2864.024 tonnes	
Average price/tonne of abatement	\$15.74	\$15.74	^c
Total revenue generated by emission abatement/year	\$127,189.277	\$45,079.744	

^a <https://www.industry.gov.au/data-and-publications/national-greenhouse-gas-inventory-quarterly-updates>

^b https://quickstats.censusdata.abs.gov.au/census_services/getproduct/census/2016/quickstat/2?opendocument

^c <http://www.cleanenergyregulator.gov.au/ERF/Pages/Auctions%20results/September%202020/Auction-September-2020.aspx>

Note: This table shows the back-of-the-envelope calculation on the total revenue generated by emission abatement per year according to the two methods of estimating the effect of goal-setting on energy usage. Column (2) uses the conventional difference-in-difference estimation results, column (3) uses the machine-learning difference-in-difference estimation results.

In Table 3.10, I perform a back of the envelop calculation on the misstatement in the revenue generated by emission abatement. I assume a monopoly registers a goal-setting emission abatement project in Victoria Australia. Gathering numbers from different sources, I find that if the government use the DiD estimates, they will overstate the revenue generated from emission abatement by \$82,109 per year. Hence, it is important to account for the household-time varying confounders that bias the estimate upward.

3.9 Conclusion

In recent years, policymakers and utility companies have a greater emphasis on energy conservation. They not only employ conventional policies such as price changes, but they also complement them with behavioural nudges. In this paper, I estimate the effect of real-time information feedback and goal-setting, a behavioural intervention that is becoming widely popular due to the rise of smart-metering and its ease of implementation.

I adopt a recently developed econometric methodology by Burlig et al. (2020) to limit the effect of selection into goal-setting, a common challenge associated with goal-setting studies using observational data. I find that real-time information feedback alone does not affect electricity usage. However, there is a scope to encourage consumers to set energy usage goals in the retail electricity context. Goal-setting reduces daily energy usage

by 0.112kWh, which is equivalent to 0.95% of the pre-treatment mean usage. A back-of-the-envelope calculation shows this reduction is equivalent to \$12 annually. Failure to control for time-varying confounders would overstate the effect of goal-setting by 64%. However, the method that I propose is useful to limit the effect of selection to the extent that it semi-parametrically controls for household-time varying confounders. Hence, the direct behavioural way consumers respond to goal setting opportunities still poses a threat to my casual identification.

As shown by [Allcott and Kessler \(2019\)](#), behavioural nudges impose a non-neglectable welfare cost on consumers in the residential electricity market. Hence, policymakers have to understand whether a specific behavioural nudge is effective in practice before any large scale implementation. My results have the following policy implications. First, the App alone has a null effect on energy usage does not necessarily imply it is welfare deteriorating due to [Allcott and Kessler \(2019\)](#)'s argument. This is because the consumer can neglect the App if it imposes a cost (e.g. cost of changing behaviour, emotional taxes etc.) unto the him, whereas the Home Energy Report in [Allcott and Kessler \(2019\)](#)'s study is unavoidable. Second, goal-setting can be effective in energy conservation, suggesting policymakers can consider to implement this behavioral intervention at a larger scale.

Chapter 4

Goals, information friction, and energy usage

4.1 Introduction

As summarized by [Thaler and Sunstein \(2009\)](#)'s seminal work, behavioral interventions have become widely adopted to influence individuals' behavior and decision making in many contexts. Policymakers and firms use these small deliberate changes to the decision environment, often without forbidding any options or significantly changing economic incentives for each option, to alter behavior by exploiting individuals' behavioral and cognitive biases¹. They are often used when conventional interventions such as price changes, taxes, and legalisations are infeasible. Goal-setting is an example of such a behavioral intervention. Even if individuals do not incur any monetary losses when they fail to achieve their goals, policymakers leverage the target's reference dependence and loss-aversion, the psychological loss incurred by individuals when they fail to achieve a goal, to generate changes in behavior ([Heath et al., 1999](#))².

Despite being a useful policy tool, information frictions can undermine the effectiveness of a goal. Being misinformed of the progress made towards achieving a goal may lead to suboptimal actions by individuals, as they may act on an incorrect belief that they have already achieved the goal. Even with the technological improvements that have occurred in recent decades, which allow individuals to monitor their goal achievement

¹Examples such as "Smart Disclosure" in the United States and the "Midata project" in the UK are innovative new tools designed to help consumers make better-informed decisions and benefit from new products and services powered by data in sectors such as health care, education, finance, energy, transportation, and telecommunications.

²It has also been shown individuals can alleviate self-control problems by setting "realistic" goals [Hsiaw \(2013\)](#). While setting the goal as a commitment device helps to alleviate self-control problem, this paper focuses on what happens *after* the goal is set.

progress with high-quality information, individuals may ignore this information due to the time and effort that it requires to access and monitor this information. For example, an individual can set a daily movement goal to motivate himself to exercise. Even if he does not incur conventional economic losses when he fails to accomplish his goal, he still exerts effort to fulfill it. Since the individual may not be able to track the number of steps he takes throughout the day, he may purchase a SmartWatch that gives him this information. However, he may not check the SmartWatch to acquire information on his progress every moment, as that may cause disruptions to his other activities. Being misinformed may lead to walk fewer steps and missing the goal. This raises the question: does the cost of acquiring information affect the effectiveness of goals?

This paper investigates how goals, information, information feedback, and reminders interact to affect individual's behavior. I develop a novel empirical framework for studying reference-dependence utility with costly information acquisition. I apply the framework to the residential electricity industry, where I examine how individuals acquire information using an App to achieve their energy goals (a specific type of reference point)³. I obtain propriety data from an electricity retailer, where they provide households with an App that gives them access to their real-time information on electricity consumption through weekly Emails and an online portal. The App also offers an option for households to set a budget as a goal for themselves. I observe households' electricity consumption data through their smart meters and their information acquisition behavior through their activity logs with the App. This novel dataset gives me an accurate measure of information acquisition, goals, and energy usage, which allows me to estimate the model.

In Chapter 3, I evaluated the App and estimated that the treatment effect of goal-setting corresponds to an annual saving of \$12. In this chapter, I focus on the goal-setters in the sample and investigate the underlying mechanisms behind this treatment effect. I specify and estimate a dynamic structural model of individual decision-making that exhibits costly information acquisition and reference dependence while allowing for biased beliefs, meaning the decision maker's belief of his progress towards goal achievement can systematically deviate from actuality. The decision-maker chooses consumption that maximizes their utility daily, but his perceived consumption may differ from actuality in the absence of immediate feedback. He may overcome this discrepancy with costly information acquisition. At the terminal period, he may experience disutility if he consumes more than the goal that he sets. I estimate this dynamic structural model to recover four behavioral parameters: (i) the disutility of violating a goal (ii) the cost of acquiring information and (iii) biasedness of households' perceived consumption compared to the actuality and (iv) the effect of a reminder. The model is also quite general, and it can be

³Researchers do not need a goal to apply my model to other environments, it is generalizable to other reference points.

applied to decision environments with reference-dependence utility (not limited to goals) and costly information acquisition.

A structural approach is appealing in this setting because it allows us to isolate different interconnected behavioral mechanisms. Isolating these channels is important, as the implication to nudges design and choice architecture are vastly different depending on which behavior is exhibited. If individuals care about their goals (reference-dependent), but the cost of acquiring information deters individuals from achieving their goals, to achieve changes in behavior, policymakers can send more reminders. On the other hand, if an individual does not care about their goals but the cost of acquiring information is low, then policymakers can provide financial incentives for goal achievement to encourage behavior change instead of sending constant reminders.

I find that households' belief of their energy consumption is unbiased without acquiring real-time information through the App. The cost of information acquisition cost reduces households' probability of acquiring information by 43%. However, since households' belief is unbiased, not being informed does not affect energy usage. Furthermore, to avoid the disutility of missing their goals, households reduce their consumption by 13%. Hence, there is little role for policymakers to nudge households into acquiring information to avoid violation of goals when households track their usage well in the absence of feedback.

4.1.1 Related literature

I contribute to the empirical literature on the effect of goals. Goals are a specific form of a reference point, as decisions are likely influenced by their relative comparisons. Reference-dependent utility is observed in various domains such as job search ([DellaVigna et al., 2017](#)), taxi drivers ([Crawford and Meng, 2011](#); [Farber, 2008](#); [Thakral and Tô, WP](#)), real effort tasks in a laboratory ([Abeler et al., 2011](#)), sports ([Allen et al., 2016](#); [Pope and Schweitzer, 2011](#)) and more. Specifically, studying goals as reference points, several studies have shown that goals are effective in many settings such as education ([Clark et al., 2020](#)), the workplace for workers' performance ([Goerg and Kube, 2012](#)) and weight-loss ([Uetake and Yang, 2018](#)).

Closest to this study, [Uetake and Yang \(2018\)](#) study the effect of goals in a calorie and weight management App using a dynamic structural model. Similar to previous studies, they assume that decision makers are fully informed of their progress with respect to their progress towards goal achievement, so they can abstract from information acquisition decisions. While this is a reasonable assumption in many settings where goals

are salient and progress is directly observable⁴, these models become less applicable in settings where the decision-maker cannot directly track their behavior, for example, the number of steps an individual makes in a morning. My model directly accounts for the decision maker's belief on the progress of goal achievement, hence creating an incentive for information acquisition as it resolves his belief uncertainty and gives him an accurate feedback on his goal achievement progress. Incorporating information acquisition is important in these models because reference dependence could be misstated if information friction is high enough to deter them from acquiring the information about their progress towards goal achievement. The implication to nudge design and choice architecture vastly differs depending on which behavior is exhibited.

Specific to the effect of goals in the residential utility setting, [Harding and Hsiaw \(2014\)](#) finds that consumers are responsive to self-set, non-binding goals. Consumers who choose realistic goals persistently save substantially more than those choosing very low or unrealistically high goals. [Agarwal et al. \(2017\)](#) confirms these descriptive statistics in a field experiment, where the authors randomly assign goals to Singaporean households to conserve water. I attempt to improve upon the current literature by incorporating how consumers acquire information using my unique dataset that records consumers' interaction with my information treatment. Furthermore, I examine how real time information programs, a commonly used nudge, interacts and complements other behavioural interventions.

My model closely relates to models of bill shocks, a phenomenon when a consumer receives a bill that is higher than what he expects. [Grubb and Osborne \(2014\)](#) and [Nevo et al. \(2016\)](#) model bill shocks in a three-part-tariff setting: a fixed fee for a usage allowance and an overage charge if consumers exceed the plan allowance. The goal in my model is functionally similar to bill shocks models' usage allowance, as they both induce dis-utility if the goal is not met or the usage exceeds the allowance. These studies ignore information acquisition decisions and assume consumers to be fully informed within the billing cycle. [Grubb \(2014\)](#) relaxes this assumption and studies both information acquisition and consumption choices in a simplified two-period model, but the author does not estimate the model. I extend the literature by providing an alternative channel to bill shocks: the lack of information acquisition or feedback within the billing cycle when consumers cannot directly track their usage.

The remainder of the paper is structured as followed. Section [4.2](#) gives a brief description of some key behaviours that I intend to model. Section [4.3](#) and [4.4](#) presents the model and the numerical DC-EGM solution method. I discuss the identification of the structural

⁴For example, it is reasonable to assume golfers to know the number of strokes they have made ([Pope and Schweitzer, 2011](#)).

parameters in section 4.5 and outline the estimation method in section 4.6. Section 4.7 presents my result. Lastly, I conclude in section 4.8.

4.2 Data

I utilize the data from the experiment that I described in Chapter 3 section 3.4. In this chapter, I focus on the 802 households (2% of the treated sample) who set a goal within the experimental period. I utilize the activity log data that records households' information acquisition behaviour and their goals, their smart-meter electricity consumption data, and their account information such as bills timing and prices per kWh. I define information acquisition as checking email, logging in the portal either through the App or the web, or clicking on an Email that directs the consumer to the portal.

FIGURE 4.1: Goal-setting households in the experiment





Figure 4.1 illustrates the sample that I use in this chapter. As I have discussed in section 3.4.2, goal-setters are different from non-goal-setters, and they are not representative of the general population. I will further discuss the implications of sample selection in my conclusion.

4.2.1 Reference dependence and rational inattention

I investigate whether engaging with the App by acquiring information helps consumers to achieve their goals. I define a consumer as engaged for a given billing cycle if he checks or clicks on an Email or logs in the portal either through the App or the web at least once within a billing cycle. I plot the kernel density of violation in goal for every bill after the household sets a goal in Figure 4.2, where positive values on the x-axis represent spending less than the goal (achieving their goals) and negative values represent spending more than the goal (falling short in goal achievement).

First, I observe that if a household never engages with the App and acquires information within a month, the distribution is centered to the left of zero, meaning households on average do not meet their goal. Second, if a household engages with the App and acquires information within a month at least once, the distribution shifts right and centered more around zero. I interpret these two facts as likely suggestive evidence of loss aversion together with rational inattention. That is, households dislike not reaching their goals. They require active information acquisition to achieve this, but there exist frictions (for example, a hassle cost of acquiring information) preventing them from doing so.

FIGURE 4.2: Kernel density and means for violation in goal



Notes: This left panel plots the kernel density of violation in goal (difference between actual bill and goal) by whether the household engage with the App during the month. I see a shift in the distribution to the right (less violation) and more centered at zero. This is confirmed by the right panel, where I plot the mean between the two distributions.

4.2.2 Within-billing cycle dynamics

I document within-billing cycle dynamics à la [Nevo et al. \(2016\)](#). From Figure 4.2, households are sensitive to their goal, and they aim to consume less than their goals. We expect households to cut down their consumption in early part of the billing cycle if they over-consumed relative to their goals, as they would re-allocate the remainder of the budget for the rest of the billing cycle. I test these implications with the following regression,

$$\begin{aligned}
 Y_{imt} = & \sum_{k=1}^4 \sum_{n=1}^5 \left(\left(\alpha_{nk} \mathbb{I} \left[pct_n \leq \frac{\sum_{\tau=1}^{t-1} kWh_{im\tau}}{G_{im}} < pct_{n+1} \right] \times \mathbb{I}[day_{mk} \geq t \geq day_{mk+1}] \right. \right. \\
 & + \alpha_k \mathbb{I}[day_{mk} \geq t \geq day_{mk+1}] + \alpha_m \mathbb{I} \left[pct_n \leq \frac{\sum_{\tau=1}^{t-1} kWh_{im\tau}}{G_{im}} < pct_{n+1} \right] \left. \right) \\
 & + \mu_i + \lambda_t + \epsilon_{imt}
 \end{aligned} \tag{4.1}$$

where Y_{imt} is the outcome of interest for household i , bill m , day t . There are two outcomes of interest: electricity consumption in kWh and information acquisition. I define information acquisition as a binary variable that equals to one if the household acquires information by either checking an Email or logging into the portal via the App. G_{im} is the goal set by household i for bill m .

$\frac{\sum_{\tau=1}^{t-1} kWh_{im\tau}}{G_{im}}$ is the proportion of the goal used up until day t . $pct_1 = 0.2$, $pct_2 = 0.4$, $pct_3 = 0.6$, $pct_4 = 0.8$, $pct_5 = 1$, $pct_6 = \infty$, so the first set of indicators, $\mathbb{1}\left[pct_n \leq \frac{\sum_{\tau=1}^{t-1} kWh_{im\tau}}{G_{im}} < pct_{n+1}\right]$, equals 1 when the proportion of the household's budget that has been used to date is between pct_n and pct_{n+1} . Similarly, I split the number of days in bill m into four equal segments. day_{mk} represents the $\left(\frac{\text{no. days in Bill}_m}{k}\right)^{th}$ day of Bill m . So the second set of indicators, $\mathbb{1}[day_{mk} \geq t \geq day_{mk+1}]$, equals 1 when day t 's observation is between day_{mk} and day_{mk+1} . Lastly, I control for household fixed effects μ_i and date fixed effects λ_t .

α_{nk} are the coefficients of interest, as they describe how households adjust their behaviours at different times of the billing cycle, depending on how well they are tracking with their goals. I present the estimated results for α_{nk} for consumption in Table 4.1 and information acquisition in Table 4.2.

TABLE 4.1: Within-bill dynamics - energy usage

	$\mathbb{1}[t \text{ in 2nd quarter}]$	$\mathbb{1}[t \text{ in 3rd quarter}]$	$\mathbb{1}[t \text{ in 4th quarter}]$
$\mathbb{1}[0 \leq \frac{C_{im(t-1)}}{G_m} < 0.4]$	-0.113	0.446	0.917***
$\mathbb{1}[0.4 \leq \frac{C_{im(t-1)}}{G_m} < 0.6]$	-5.450***	-6.161***	-6.028***
$\mathbb{1}[0.6 \leq \frac{C_{im(t-1)}}{G_m} < 0.8]$	-5.418***	-6.902***	-7.304***
$\mathbb{1}[0.8 \leq \frac{C_{im(t-1)}}{G_m} < 1]$	-2.460***	-4.757***	-5.803***
$\mathbb{1}[1 \leq \frac{C_{im(t-1)}}{G_m}]$	3.603***	2.430***	1.502***

*** p<0.01, ** p<0.05, * p<0.1

Notes: This table presents the OLS estimation result of α_{nk} for Equation (4.1), with energy usage as the outcome. $C_{im(t-1)} = \sum_{\tau=1}^{t-1} kWh_{im\tau}$ is the aggregate consumption used up until the previous day. The 1st quarter of the billing cycle is the omitted category. The regression includes day of the cycle fixed effect, day fixed effect and household fixed effect.

TABLE 4.2: Within-bill dynamics - information acquisition

	$\mathbb{1}[t \text{ in 2nd quarter}]$	$\mathbb{1}[t \text{ in 3rd quarter}]$	$\mathbb{1}[t \text{ in 4th quarter}]$
$\mathbb{1}[0 \leq \frac{C_{im(t-1)}}{G_m} < 0.4]$	0.026*	0.028	0.026
$\mathbb{1}[0.4 \leq \frac{C_{im(t-1)}}{G_m} < 0.6]$	0.027	0.029*	0.036
$\mathbb{1}[0.6 \leq \frac{C_{im(t-1)}}{G_m} < 0.8]$	0.033	0.038	0.037
$\mathbb{1}[0.8 \leq \frac{C_{im(t-1)}}{G_m} < 1]$	-0.027	-0.023	-0.024
$\mathbb{1}[1 \leq \frac{C_{im(t-1)}}{G_m}]$	0.005	0.022	0.017

*** p<0.01, ** p<0.05, * p<0.1

Notes: This table presents the OLS estimation result of α_{nk} for Equation (4.1), with information acquisition as the outcome. $C_{im(t-1)} = \sum_{\tau=1}^{t-1} kWh_{im\tau}$ is the aggregate consumption used up until the previous day. The 1st quarter of the billing cycle is the omitted category. The regression includes day of the cycle fixed effects, day fixed effects and household fixed effects.

From Table 4.1, we observe that households adjust their usage in accordance to whether they think they will meet their goals. First, when households are unlikely to fall short in meeting their goal (represented by the first row), they do not adjust their consumption. However, if they are as they get closer to violating their goals (represented by rows three and four), they reduce their consumption, and the adjustment is larger later in the billing cycle (columns two and three). Conversely, if they have already missed their goals (represented by the final row), instead of exerting the effort to reduce energy usage, they would consume more. I do not observe these patterns for information acquisition from Table 4.2.

4.2.3 Effect of Email reminders

While I do not find patterns in the timing of information acquisition with respect to the progress of goal achievement, I investigate whether reminders nudge households into engaging with the App by acquiring information. I run the following regression

$$\text{AcquireInformation}_{it} = \alpha_i + \tau_t + \beta_1 \text{EmailReminder}_{it} + \varepsilon_{it} \quad (4.2)$$

where $\text{AcquireInformation}_{it}$ equals one if household i checks the Email or logs the portal on day t , and $\text{EmailReminder}_{it}$ equals one if I send household i an Email reminder on day t . α_i and τ_t are household and date fixed effects respectively. I cluster the standard errors at the household level. β_1 has a casual interpretation, representing the change in the likelihood of information acquisition when they are being reminded.

TABLE 4.3: The effect of an Email reminder on information acquisition

	(1)	(2)
	Acquires information	Acquires information
EmailReminder	0.359*** (0.014)	0.333*** (0.015)
Constant	0.060*** (0.005)	0.064*** (0.002)
Observations	259,414	259,406
R-squared	0.155	0.283
Household fixed effects	N	Y
Date fixed effects	N	Y

Robust standard errors in parentheses

*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Notes: This table present the estimation result for Equation (4.2). The outcome variable is a binary variable that indicates whether the consumer acquires information by checking an Email or logging in to the portal. The variable of interest is a binary variable that indicates whether the consumer receives an Email from the App. I observe consumers have a 33% higher probability of acquiring information on a day they receive an Email reminder.

As Table 4.3 shows, households are 33% more likely to acquire information on a day where they receive an Email reminder. This is consistent with the previous studies on reminder nudges (Calzolari and Nardotto, 2017; Karlan et al., 2016, for example), as they are designed to bring a particular decision to the recipient's attention.

4.3 Model

As motivated by the descriptive statistics presented in the earlier section, I aim to model a dynamic decision problem that incorporates (i) reference dependence, (ii) costly information acquisition and (iii) effective reminders. Appendix C.1 provides a comprehensive list of variables that I use in the model and their descriptions.

4.3.1 Model setup

The forward-looking decision-maker (DM) solves a maximisation problem over a billing cycle that lasts for T days. A decision period t is a day. When the DM enters period t , he

has a belief of the *remainder of the goal*, G_{t-1} (i.e. the goal minus the belief of consumption up to $t - 1$). The belief is fully characteristics by the two state variables $\tilde{G}_{t-1}$ and $\tilde{V}G_{t-1}$ ⁵, the mean and the variance respectively. I assume this to be normally distributed. I interpret this as the DM's *prior* belief in period t . That is, $G_{t-1} \sim N(\tilde{G}_{t-1}, \tilde{V}G_{t-1})$, with $\tilde{G}_0$ given and he knows this for certain ($\tilde{V}G_0 = 0$) at the start of the billing cycle⁶.

Within each day, a reminder r_t is first realised. This is a binary variable that represents whether the DM receives a reminder in this period. Whether the DM receives a reminder in future periods r_{t+k} is stochastic, but the DM knows the arrival of reminders in future periods follows a Bernoulli distribution with probability q . The DM also experiences a information acquisition cost shock $\varepsilon_t(i_t) \sim EV1(\sigma)$.

Upon observing these two shocks, the DM makes his consumption (in kWh) choice c_t and information acquisition choice i_t . Consumption is continuous and information acquisition is discrete. Electricity consumption generates flow utility. I assume a constant relative risk aversion (CRRA) form of consumption utility. The DM's indirect utility with the curvature parameter of β and price per kWh of p :

$$u(c_t) = \frac{c_t^{1-\beta} - 1}{1-\beta} - pc_t \quad (4.3)$$

If the DM acquires information, his belief will be updated such that the mean will be perfectly aligned with the actual goal remainder with no uncertainty in $t + 1$. Such action is costly, and it incurs a cost of κ . I interpret this as the opportunity cost of time, as the DM has to scroll through Emails or spend time logging into the App.

If the DM acquires information in a period where he receives a reminder ($r_t = 1$), the information acquisition cost is reduced by ν . I remain agnostic about the sign of ν . I hypothesize it to be positive, as a reminder like an App notification reduces the cost of acquiring information. The DM may have to login and scroll through their Emails to find this information in the absence of the reminder. However, there can be an additional cost to information acquisition if the reminder imposes a nuisance cost as suggested by [Damgaard and Gravert \(2018\)](#). The information cost function can be written as:

$$C(i_t, r_t) = i_t(\kappa - r_t\nu) + \varepsilon_t(i_t) \quad (4.4)$$

⁵I want to note that this can be any reference point in general, but I will use *goal* for the rest of my exposition as it suits this context.

⁶In [Harding and Hsiaw \(2014\)](#)'s model, a present-biased consumer with reference-dependent utility encounters a two-stage optimization problem. In the first stage, they decide whether to enrol in an energy saver program that assigns him a goal. In the second stage, they choose monthly energy consumption (with a goal if they enrolled in the first stage). In their model, present-biasedness is necessary to explain why the consumer would sign up for the program in the first place. By treating the goal $\tilde{G}_0$ as exogenous, my model abstract from their first stage decision and focus on the second stage.

If the DM does not acquire information, then his updated belief will be the prior less this period consumption with noise, where the noise is $\xi_t \sim N(\mu_{\xi}, \sigma_{\xi})$, and the variance of belief will increase due to more uncertainty from the noise⁷.

I want to note that the actual amount of the goal remainder is realized at the end of the period; that is, *after* the information acquisition decision made in each period. Hence, ex-ante to the reveal of the information, the DM has to make an inference on what he thinks the mean of the bill is if he acquires information. If I abstract from other behavioral biases such as motivated beliefs/memory, it is reasonable to assume that the mean of the belief of the actual bill if the DM acquires information, prior to acquiring this information, is his prior, less this period's consumption.

Hence, the mean of the belief $\tilde{G}_t$ (ex-ante) evolves according to

$$\tilde{G}_t = i_t(\tilde{G}_{t-1} - pc_t) + (1 - i_t)(\tilde{G}_{t-1} - p(c_t + \xi_t)) \quad (4.5)$$

and ex-post, if the DM acquires information, his budget remainder will be his initial budget less aggregate actual consumption ($\tilde{G}_0 - p \sum_{j=1}^t c_j$). The variance of the DM's belief $\tilde{V}G_t$ evolves according to

$$\tilde{V}G_t = (1 - i_t)(\tilde{V}G_{t-1} + \sigma_{\xi}^2) \quad (4.6)$$

At the end of the billing cycle, the DM's bill is realized. I model the DM to be reference dependent, and the reference point being the goal that he is given at the start $\tilde{G}_0$. I allow the DM to be loss averse, meaning that he experiences more utility loss in the circumstance that he consumes more than he goal he sets ($\tilde{G}_T < 0$) than the utility gain he experiences when he consumes less than he goal he sets ($\tilde{G}_T > 0$) for the same amount of absolute over/under consumption. That is

$$GL(\tilde{G}_T, \tilde{V}G_T) = \begin{cases} \eta G_T & \text{if } G_T > 0 \\ \eta \lambda G_T & \text{otherwise} \end{cases} \quad (4.7)$$

$$\text{where } G_T \sim N(\tilde{G}_T, \tilde{V}G_T) \quad (4.8)$$

⁷In [Grubb \(2014\)](#)'s simple two-period model, similar to my model, if the consumer acquires information, he can condition his second-period consumption decision on whether he consumed in the first period. When the consumer does not acquire information, their model assumes the consumer can only condition his second-period consumption strategy on his first-period consumption strategy. In contrast, I model consumers to have a perceived consumption $c_t + \xi_t$ as a signal to their actual consumption c_t . And since ξ_t is i.i.d Gaussian, the variance of their belief is $\tilde{V}G_{t-1}$ (variance of prior) + σ_{ξ}^2 (additional variance from noise). This approach is supported by [Attari et al. \(2014\)](#), as they suggest that households have difficulties judging the actual amount of electricity usage.

where η represents the sensitivity to the reference point and λ represents loss aversion. The DM maximizes the expected sum of the flow utility with the terminal period's utility. The maximization problem can be written as⁸

$$\begin{aligned}
 & \max_{\{c_t, i_t\}_{t=1}^T} E \sum_{t=1}^T \left[\frac{c_t^{1-\beta} - 1}{1-\beta} - pc_t - i_t(\kappa - r_t v) + \varepsilon(i_t) + GL(G_t) \right] \\
 \text{s.t. } & i_t \in \{0, 1\}, c_t \in \mathbb{R}^+ \\
 & G_t \sim N(\tilde{G}_t, \tilde{V}G_t) \\
 & \tilde{G}_t = i_t(G_{t-1} - pc_t) + (1 - i_t)(\tilde{G}_{t-1} - p(c_t + \xi_t)) \\
 & \tilde{V}G_t = (1 - i_t)(\tilde{V}G_{t-1} + \sigma_\xi^2) \\
 & r_t \sim \text{Bernoulli}(q) \\
 & \varepsilon_t(i_t) \stackrel{iid}{\sim} EV1 \\
 & \xi_t \stackrel{iid}{\sim} N(\mu_\xi, \sigma_\xi^2) \\
 & \tilde{G}_0, \tilde{V}G_0 \text{ given}
 \end{aligned}$$

where gain loss utility $GL(\cdot)$ is defined as:

$$GL(x) = \begin{cases} \eta x & \text{if } x > 0 \\ \eta \lambda x & \text{otherwise} \end{cases}$$

4.3.2 Solving the model

The maximization problem can be characterized by the following Bellman equations:

For $t \leq T$:

$$V_t(\tilde{G}_{t-1}, \tilde{V}G_{t-1}, r_t, \varepsilon_t) = \max_{c_t \geq 0, i_t \in \{0, 1\}} [u(c_t) - C(i_t, r_t) + \varepsilon_t(i_t) + E(V_{t+1}(\tilde{G}_t, \tilde{V}G_t \mid \tilde{G}_{t-1}, \tilde{V}G_{t-1}, c_t, i_t))]$$

For $t = T + 1$:

$$V_t(\tilde{G}_{t-1}, \tilde{V}G_{t-1}, r_t, \varepsilon_t) = GL(\tilde{G}_{t-1}, \tilde{V}G_{t-1} \mid \tilde{G}_{t-1}, \tilde{V}G_{t-1}, c_t, i_t) \tag{4.9}$$

⁸Similar to [Nevo et al. \(2016\)](#), since these are daily consumption decisions over a short horizon, I ignore inter-temporal discounting. However, this also means I omit the present-bias feature of the decision problem, that I heavily emphasized in Chapter 3. As [Harding and Hsiaw \(2014\)](#)'s theoretical model suggests, present-biasedness is necessary to explain why consumers will set a goal at the first place, as they use the goal as a commitment device to curb their overconsumption due to their time inconsistency. I omit this feature due to the identification concerns. Since this is out of the scope of this essay, I refer interested readers to Appendix C.2.

I can rewrite (4.9) as

$$V_t(\tilde{G}_{t-1}, \tilde{V}G_{t-1}, r_t, \varepsilon_t) = \max_{i_t \in \{0,1\}} \begin{cases} w_t(\tilde{G}_{t-1}, \tilde{V}G_{t-1}, r_t, 0) + \sigma \varepsilon_t(0) \\ w_t(\tilde{G}_{t-1}, \tilde{V}G_{t-1}, r_t, 1) + \sigma \varepsilon_t(1) \end{cases} \quad (4.10)$$

where $w_t(\cdot, i_t)$ are choice specific value functions:

$$\begin{aligned} w_t(\tilde{G}_{t-1}, \tilde{V}G_{t-1}, r_t, 0) &= \max_{c_t} [u(c_t) + EV_{t+1}(\tilde{G}_{t-1} - p(c_t + \xi_t), \tilde{V}G_{t-1} + \sigma \xi_t, r_{t+1}, \varepsilon_{t+1})] \\ w_t(\tilde{G}_{t-1}, \tilde{V}G_{t-1}, r_t, 1) &= \max_{c_t} [u(c_t) - (\kappa - r_t \nu) + EV_{t+1}(\tilde{G}_{t-1} - pc_t, 0, r_{t+1}, \varepsilon_{t+1})] \end{aligned} \quad (4.11)$$

where the expectation is respect to r_{t+1} , ε_{t+1} and ξ_t (realized at the end of the period). Using the distributional assumption of r and ε , $EV_{t+1}(\cdot)$ has the logsum representation (McFadden, 1973):

$$\begin{aligned} EV_{t+1}(\tilde{G}_t, \tilde{V}G_t) &= E[q(\sigma \log[\exp(w_{t+1}(\tilde{G}_t, \tilde{V}G_t, 1, 0)/\sigma) + \exp(w_{t+1}(\tilde{G}_t, \tilde{V}G_t, 1, 1)/\sigma)]) \\ &\quad + (1 - q)(\sigma \log[\exp(w_{t+1}(\tilde{G}_t, \tilde{V}G_t, 0, 0)/\sigma) + \exp(w_{t+1}(\tilde{G}_t, \tilde{V}G_t, 0, 1)/\sigma)])] \end{aligned} \quad (4.12)$$

where the expectation now is with respect to ξ_t . With a special case when $t = T$, $EV_{T+1}(\cdot) = EGL(\cdot)$, using the functional form of $EGL(\cdot)$, it can be written as

$$EGL(\tilde{G}_t, \tilde{V}G_t) = \int_{\xi} \frac{1}{\sqrt{2\pi(\tilde{V}G_t)}} \int_G \exp\left(-\frac{G - (\tilde{G}_t - p(c_T + \xi))}{2(\tilde{V}G_t)^2}\right) ((G > 0)\eta G + (G < 0)\eta \lambda G) dG dF(\xi) \quad (4.13)$$

I can finally rewrite the conditional value functions with this general representation

$$w_t(\tilde{G}_{t-1}, \tilde{V}G_{t-1}, r_t, i_t) = \max_{c_t} [u(c_t) - i_t(\kappa - r_t \nu) + E(LS(w_{t+1}(\tilde{G}_t, \tilde{V}G_t, i_{t+1}) | \tilde{G}_{t-1}, \tilde{V}G_{t-1}, c_t, i_t))] \quad (4.14)$$

where $LS(\cdot)$ is the logsum equation. The solution to (4.9)-(4.14) gives the optimal decision rules:

$$\begin{aligned} c_t^*(\tilde{G}_{t-1}, \tilde{V}G_{t-1}, r_t) &: [\tilde{G}_{t-1}, \tilde{V}G_{t-1}, r_t] \rightarrow \mathbb{R}^+ \\ i_t^*(\tilde{G}_{t-1}, \tilde{V}G_{t-1}, r_t) &: [\tilde{G}_{t-1}, \tilde{V}G_{t-1}, r_t] \rightarrow [P_t(n=0, i=0), P_t(n=0, i=1), P_t(n=1, i=0), P_t(n=1, i=1)] \end{aligned} \quad (4.15)$$

where $P_t(r, i)$ is the probability of ignoring information ($i = 0$) or acquiring information ($i = 1$), contingent on whether the DM receives a reminder ($r = 1$) or not ($r = 0$) at period

t , given by

$$\begin{aligned}
P_t(r_t = 0, i_t = 0) &= \frac{\exp(w_t(\tilde{G}_{t-1}, \tilde{V}_{G_{t-1}}, 0, 0)/\sigma)}{\exp(w_t(\tilde{G}_{t-1}, \tilde{V}_{G_{t-1}}, 0, 0)/\sigma) + \exp(w_t(\tilde{G}_{t-1}, \tilde{V}_{G_{t-1}}, 0, 1)/\sigma)} \\
P_t(r_t = 0, i_t = 1) &= \frac{\exp(w_t(\tilde{G}_{t-1}, \tilde{V}_{G_{t-1}}, 0, 1)/\sigma)}{\exp(w_t(\tilde{G}_{t-1}, \tilde{V}_{G_{t-1}}, 0, 1)/\sigma) + \exp(w_t(\tilde{G}_{t-1}, \tilde{V}_{G_{t-1}}, 1, 0)/\sigma)} \\
P_t(r_t = 1, i_t = 0) &= \frac{\exp(w_t(\tilde{G}_{t-1}, \tilde{V}_{G_{t-1}}, 1, 0)/\sigma)}{\exp(w_t(\tilde{G}_{t-1}, \tilde{V}_{G_{t-1}}, 1, 0)/\sigma) + \exp(w_t(\tilde{G}_{t-1}, \tilde{V}_{G_{t-1}}, 1, 1)/\sigma)} \\
P_t(r_t = 1, i_t = 1) &= \frac{\exp(w_t(\tilde{G}_{t-1}, \tilde{V}_{G_{t-1}}, 1, 1)/\sigma)}{\exp(w_t(\tilde{G}_{t-1}, \tilde{V}_{G_{t-1}}, 1, 1)/\sigma) + \exp(w_t(\tilde{G}_{t-1}, \tilde{V}_{G_{t-1}}, 1, 1)/\sigma)}
\end{aligned} \tag{4.16}$$

4.4 Numerical solution

I solve the model numerically using the DC-EGM algorithm developed by [Iskhakov et al. \(2017\)](#). This method is appealing compared to the traditional value function iteration (VFI) method as it avoids costly root-finding operations that would be required to find the optimal consumption at each point of the exogenous state grid. Furthermore, since this maximization problem involves one discrete choice and one continuous choice variable, the first order condition of the Euler equation is generally non-concave, so the solution cannot rely on the first order condition. The DC-EGM algorithm automatically detects the region of non-concavity and adjusts for them accordingly. It has been shown that this method is faster and more accurate compared to the VFI.

To utilize the DC-EGM algorithm, I have to first derive the Euler equation. Note that, conditional on the discrete information acquisition choice i_t , optimal consumption choices are characterized by the Bellman Equations (4.14). The first order condition for consumption is given by

$$\begin{aligned}
u'(c_t, i_t) &= E \left[q \sum_{i_{t+1} \in \{0,1\}} P_{t+1}(r_{t+1} = 1, i_{t+1}) \frac{\partial w_{t+1}(\tilde{G}_t, \tilde{V}_{G_t}, r_{t+1} = 1, i_{t+1})}{\partial \tilde{G}_t} p \mid \tilde{G}_{t-1}, \tilde{V}_{G_{t-1}} \right. \\
&\quad \left. + (1 - q) \sum_{i_{t+1} \in \{0,1\}} P_{t+1}(r_{t+1} = 0, i_{t+1}) \frac{\partial w_{t+1}(\tilde{G}_t, \tilde{V}_{G_t}, r_{t+1} = 0, i_{t+1})}{\partial \tilde{G}_t} p \mid \tilde{G}_{t-1}, \tilde{V}_{G_{t-1}} \right]
\end{aligned} \tag{4.17}$$

where $P_{t+1}(r_{t+1}, i_{t+1})$ is defined in (4.16). From (4.14), using the envelope theorem,

$$\frac{\partial w_t(\tilde{G}_{t-1}, \tilde{V}_{G_{t-1}}, i_t)}{\partial \tilde{G}_{t-1}} = E \left[\frac{\partial w_{t+1}(\tilde{G}_t, \tilde{V}_{G_t}, r_{t+1}, i_{t+1})}{\partial \tilde{G}_t} \frac{\partial \tilde{G}_t}{\partial \tilde{G}_{t-1}} \mid \tilde{G}_{t-1}, \tilde{V}_{G_{t-1}} \right] \tag{4.18}$$

since $\partial \tilde{G}_t / \partial \tilde{G}_{t-1} = 1 = -p \partial \tilde{G}_t / \partial c_t$, it holds for $i_t = 0, 1$ and $t = 1, \dots, T-1$,

$$u'(c_t, i_t) = \frac{\partial w_t(\tilde{G}_{t-1}, \tilde{V}G_{t-1}, r_t, i_t)}{\partial \tilde{G}_{t-1}} p \quad (4.19)$$

Hence, in $t+1$, it holds too:

$$u'(c_{t+1}, i_{t+1}) = \frac{\partial w_{t+1}(\tilde{G}_t, \tilde{V}G_t, r_{t+1}, i_{t+1})}{\partial \tilde{G}_t} p \quad (4.20)$$

Substitute (4.20) into (4.17),

$$\begin{aligned} u'(c_t, i_t) = E \Bigg[& q \sum_{i_{t+1} \in \{0,1\}} P_{t+1}(r_{t+1} = 1, i_{t+1}) u'(c_{t+1}, i_{t+1}) \\ & + (1-q) \sum_{i_{t+1} \in \{0,1\}} P_{t+1}(r_{t+1} = 0, i_{t+1}) u'(c_{t+1}, i_{t+1}) \mid \tilde{G}_{t-1}, \tilde{V}G_{t-1} \Bigg] \end{aligned} \quad (4.21)$$

I want to stress that the Euler equation is *conditional* on information acquisition choice i_t ⁹.

Specifically, for the decision to not acquire information $i_t = 0$,

$$\begin{aligned} u'(c_t, 0) = \int [& q P_{t+1}(1, 0 \mid \tilde{G}_{t-1} - p(c + \xi), \tilde{V}G_{t-1} + \sigma_\xi) u'(c_{t+1}, 0 \mid \tilde{G}_{t-1} - p(c + \xi), \tilde{V}G_{t-1} + \sigma_\xi) \\ & + q P_{t+1}(1, 1 \mid \tilde{G}_{t-1} - p(c + \xi), \tilde{V}G_{t-1} + \sigma_\xi) u'(c_{t+1}, 1 \mid \tilde{G}_{t-1} - p(c + \xi), \tilde{V}G_{t-1} + \sigma_\xi) \\ & + (1-q) P_{t+1}(0, 0 \mid \tilde{G}_{t-1} - p(c + \xi), \tilde{V}G_{t-1} + \sigma_\xi) u'(c_{t+1}, 0 \mid \tilde{G}_{t-1} - p(c + \xi), \tilde{V}G_{t-1} + \sigma_\xi) \\ & + (1-q) P_{t+1}(0, 1 \mid \tilde{G}_{t-1} - p(c + \xi), \tilde{V}G_{t-1} + \sigma_\xi) u'(c_{t+1}, 1 \mid \tilde{G}_{t-1} - p(c + \xi), \tilde{V}G_{t-1} + \sigma_\xi)] dF(\xi) \end{aligned} \quad (4.22)$$

and for the decision to acquire information $i_t = 1$,

$$\begin{aligned} u'(c_t, 1) = & [q P_{t+1}(1, 0 \mid \tilde{G}_{t-1} - pc, 0) u'(c_{t+1}, 0 \mid \tilde{G}_{t-1} - pc, 0) \\ & + q P_{t+1}(1, 1 \mid \tilde{G}_{t-1} - pc, 0) u'(c_{t+1}, 1 \mid \tilde{G}_{t-1} - pc, 0) \\ & + (1-q) P_{t+1}(0, 0 \mid \tilde{G}_{t-1} - pc, 0) u'(c_{t+1}, 0 \mid \tilde{G}_{t-1} - pc, 0) \\ & + (1-q) P_{t+1}(0, 1 \mid \tilde{G}_{t-1} - pc, 0) u'(c_{t+1}, 1 \mid \tilde{G}_{t-1} - pc, 0) \end{aligned} \quad (4.23)$$

The four lines represents the marginal utility next period weighted by the probability of: (i) being reminded and not acquiring information next period (ii) being reminded and acquiring information next period, (iii) not being reminded and not acquiring information next period and (iv) not being reminded and not acquiring next period next period. Given the next period choice-specific consumption functions $c_{t+1}(\tilde{G}_t, \tilde{V}G_t, 0)$ and $c_{t+1}(\tilde{G}_t, \tilde{V}G_t, 1)$, I solve (4.22) and (4.23) using backward induction. Starting from period T , the DM's perception of the budget remainder $G_{T-1} \sim N(\tilde{G}_{T-1}, \tilde{V}G_{T-1})$ the DM solves

⁹It can easily be shown that it does not depend on whether the DM receives a reminder r_t . Intuitively, r_t does not enter the first order condition with respect to c_t .

for the following choice-specific consumption:

$$c_T^*(\tilde{G}_{T-1}, \tilde{V}G_{T-1}, 0) = \operatorname{argmax}_{c_T} \left[\frac{c_T^{1-\beta} - 1}{1-\beta} - pc_T + EGL(\tilde{G}_{T-1} - p(c_T + \xi_T), \tilde{V}G_{T-1} + \sigma_\xi) \right] \quad (4.24)$$

$$c_T^*(\tilde{G}_{T-1}, \tilde{V}G_{T-1}, 1) = \operatorname{argmax}_{c_T} \left[\frac{c_T^{1-\beta} - 1}{1-\beta} - pc_T - (\kappa - r_T v) + GL(\tilde{G}_{T-1} - pc_T, 0) \right] \quad (4.25)$$

with the probability of acquiring information of

$$P_T(r_T = 1, i_T = 1) = \frac{\exp(w_T(\tilde{G}_{T-1}, \tilde{V}G_{T-1}, 1, 1)/\sigma)}{\exp(w_T(\tilde{G}_{T-1}, \tilde{V}G_{T-1}, 1, 0)/\sigma) + \exp(w_T(\tilde{G}_{T-1}, \tilde{V}G_{T-1}, 1, 1)/\sigma)}$$

$$P_T(r_T = 0, i_T = 1) = \frac{\exp(w_T(\tilde{G}_{T-1}, \tilde{V}G_{T-1}, 0, 1)/\sigma)}{\exp(w_T(\tilde{G}_{T-1}, \tilde{V}G_{T-1}, 0, 0)/\sigma) + \exp(w_T(\tilde{G}_{T-1}, \tilde{V}G_{T-1}, 0, 1)/\sigma)} \quad (4.26)$$

with the choice specific value function defined as

$$w_T(\tilde{G}_{T-1}, \tilde{V}G_{T-1}, r_T, 0) = \left[\frac{c_T^{*1-\beta} - 1}{1-\beta} - pc_T^*(\tilde{G}_{T-1}, \tilde{V}G_{T-1}, 0) + EGL(\tilde{G}_{T-1} - p(c_T^* + \xi_T), \tilde{V}G_{T-1} + \sigma_\xi) \right] \quad (4.27)$$

$$w_T(\tilde{G}_{T-1}, \tilde{V}G_{T-1}, r_T, 1) = \left[\frac{c_T^{*1-\beta} - 1}{1-\beta} - pc_T^*(\tilde{G}_{T-1}, \tilde{V}G_{T-1}, 1) - (\kappa - r_T v) + GL(\tilde{G}_{T-1} - pc_T^*, 0) \right] \quad (4.28)$$

For $t < T$, I apply backward induction. For each period t , I take the already computed policy function $c_{t+1}^*(\tilde{G}_{t+1}, \tilde{V}G_{t+1})$ and choice specific value functions $w_{t+1}(\tilde{G}_t, \tilde{V}G_t, i_{t+1})$ as given. First I define $A_t = \tilde{G}_{t-1} - pc_t$ the *exogenous grid*, the sufficient statistics that carries all the information the goal remainder and consumption in the period¹⁰. Denote $A = \{A_1, \dots, A_J\}$ the monotonic exogenous grid over the *actual goal remainder*, where A_1 can be negative¹¹.

For each point A_j in the exogenous grid, I calculate the right hand side (RHS) of equation (4.21) by integrating over the *belief error* ξ_t (realized after information acquisition choice is made) using Gaussian quadrature. For each quadrature point representing a possible value of ξ_t , I compute (i) next period's belief variance state $\tilde{V}G_t$ according to (4.6), (ii) next period's belief mean state $\tilde{G}_t$ according to (4.5) using the exogenous grid A_t ¹²

¹⁰Note that $\tilde{G}_t$ only depends on $\tilde{G}_{t-1}$ and c_t only through $A_t = \tilde{G}_{t-1} - pc_t$, hence A_t serves as a sufficient statistics for period t .

¹¹I exogenously vary that in my numerical solution to check for robustness

¹²I want to stress this is as this is the innovation of the EGM. $\tilde{G}_t$ only depends $\tilde{G}_{t-1}$ and c_t through A_t .

and (iii) The optimal level of consumption in $t + 1$ using the conditional policy function $c_{t+1}(\tilde{G}_t, \tilde{V}G_t, i_{t+1})$ from backward induction.

With (i), (ii), (iii), I can compute the *RHS* of the Euler equation according to (4.21). Given the CRRA specification of utility in (4.3), Equation (4.21) is analytically invertible for the current period level of optimal consumption c_t for every A_t , it avoids root-finding for each level of the state variable, ensuring the speed for the numerical solution. I compute current period optimal consumption (condition on whether to acquire information) with $c_t(\tilde{G}_{t-1}, \tilde{V}G_{t-1}, i_t) = (u')^{-1}(RHS)$, the inverse of the marginal utility of the *RHS*. It should be noted that at this point, the policy function is defined over the exogenous grid A .

With the current period consumption policy function, I can invert back for the endogenous grid, namely $\tilde{G}_j = A_j + pc_t(\tilde{G}_{t-1}, \tilde{V}G_{t-1}, i_t)$. However, as demonstrated in Iskhakov et al. (2017), the endogenous grid $\tilde{G}_t$ may not necessarily be monotone in the exogenous grid A_t . Using the upper envelope algorithm from Iskhakov et al. (2017), I eliminate sub-optimal points corresponding to the Euler equation solution which do not correspond to the highest values of $w_t(\tilde{G}_{t-1}, \tilde{V}G_{t-1}, i_t)$.

I backward induct according to the aforementioned steps until the starting period. The advantage of this algorithm is that it avoids root-finding operation for every period except for the terminal period and uses the analytical solution of the Euler equation to compute the optimal consumption policy. Unlike the conventional VFI, where it requires computational costly root-finding for every period, and uses approximations (defined by the fineness of the grid) to compute the optimal consumption policy.

4.5 Identification

Due to the non-linearity of the parameters, it is difficult to precisely define what variation in the data identifies each parameter individually. Nevertheless, I provide a brief discussion on how each parameter can be identified in this section.

κ represents the information acquisition cost. A higher κ provides less incentive for the DM to acquire information as it decreases the utility more when he does so. Hence, we would expect the probability of acquiring information to be lower. Therefore, holding all else constant, κ can be identified by the probability of information acquisition (regardless of whether the DM receives a reminder).

ν represents the reduction in the information acquisition cost when the DM acquires information in a period where he receives a reminder. This benefit can only be realized in

periods when the DM receives a reminder ($r_t = 1$), and he decides to acquire information ($i_t = 1$). A higher ν provides the DM a higher incentive to acquire information when the DM receives a reminder, hence we would expect the probability of information acquisition to be higher in such instances. Therefore, holding all else constant, ν can be identified by comparing the probability of information acquisition with and without a reminder, i.e. $corr(i_t, r_t)$.

η , λ represents sensitivity to gain loss and loss aversion respectively. Suppose $\eta = 0$, that means the DM is completely insensitive to their gain-loss utility, so holding all other factors constant, they will not decrease their consumption in the final period when $\tilde{G}_{T-1}$ is close to zero to achieve a higher gain-loss utility. Therefore, η can be identified by comparing consumption when the DM has goal reminder that is close to 0 ($\tilde{G}_{T-1} = 0$) to consumption with high level of goal remainder (high $\tilde{G}_{T-1}$). Likewise, λ only affects the DM when $\tilde{G}_T$ is negative, hence it can be identified by the asymmetric consumption response in the final period between DM's with a positive and negative $\tilde{G}_{T-1}$. The idea is similar to [Uetake and Yang \(2018\)](#) and [Saez \(2010\)](#) which identifies the labor elasticity with respect to income tax using the bunching of income at the kink in the income tax schedule. Intuitively, if there is no reference dependence in utility, I do not observe the bunching at the kink.

β represents risk aversion, a higher β corresponds to a higher decreasing marginal utility of consumption. To illustrate identification of β from λ and η , consider a one period model. The choice specific value function for $i_1 = 1$ in the final period is from (4.28)

$$\begin{aligned} w_1(\tilde{G}_0, \tilde{V}G_0, r_1, 1) &= \max_{c_1} \left[\frac{c_1^{1-\beta} - 1}{1-\beta} - pc_1 - (\kappa - r_T\nu) + GL(\tilde{G}_0 - pc_1, 0) \right] \\ &= \max_{c_1} \left[\frac{c_1^{1-\beta} - 1}{1-\beta} - pc_1 - (\kappa - r_T\nu) + \eta\lambda\mathbb{1}(\tilde{G}_0 < pc_1)(\tilde{G}_0 - pc_1) \right. \\ &\quad \left. + \eta\mathbb{1}(\tilde{G}_0 > pc_1)(\tilde{G}_0 - pc_1) \right] \end{aligned}$$

β can be identified with the variation in p and c_1 across individuals. While the identification of the reference dependence parameters η and λ also depends on c_1 and p , as I have discussed previously, they rely on the variation in $\tilde{G}_0$ too. Furthermore, since the gain loss utility $GL(\cdot)$ are piecewise linear, and the CRRA utility function is non-linear, the functional form helps to separately identify β from λ and η .

μ_{ξ} represents the mean of the belief error. $\mu_{\xi} = 0$, happens when the DM has an unbiased belief, meaning that his actual consumption (c_t) and perceived consumption ($c_t + \xi_t$) is on average the same. $\mu_{\xi} > 0 (< 0)$ are situations when the perceived consumption exceeds (be less than) the actual consumption. Suppose $\mu_{\xi} = 0$ (i.e. unbiased belief), then on average, the DM's belief ($\tilde{G}_t$) aligns with actuality. Holding other constant, we would

expect consumption after information acquisition not to differ sharply. Suppose $\mu_{\xi} < 0$ (i.e, perceived consumption less than actual consumption), then as time progress, the DM's belief of the goal remainder ($\tilde{G}_t$) is more than actuality and the gap widens. Once the DM acquires information, his belief of the goal remainder decreases sharply when it aligns with actuality in the following period. Hence, the consumption changes after information acquisition relative to the prior help to identify μ_{ξ} .

σ_{ξ} represents the variance of the belief error. A higher variance means a less precise error, implying the DM's belief state fluctuates due to the volatility of ξ_t . We would expect the DM's consumption to be more volatile due to the volatility his belief state $\tilde{G}_t$. Hence, the volatility of consumption for a given bill helps to identify σ_{ξ} . I also want to note that, given the DM is loss aversion ($\lambda > 1$), σ_{ξ} provides the incentive for the individual to acquire information. To illustrate this, suppose $\tilde{G}_T = 0$, that is, the DM believes he just achieves the goal in this billing cycle. However, if he does not acquire information in the final period, he is uncertain about this, as his belief is $G_T \sim N(0, \tilde{V}G_T)$, where $\tilde{V}G_T$ is a function of σ_{ξ} , as the uncertainty accumulates since his last instance of information acquisition. While on average he is correct that he achieves his goal, due to loss aversion, the expected utility of gain-loss given by (4.13) is negative¹³. This incentivizes the DM the acquire information to resolve this uncertainty, so he is certain that he did not violate his goal.

4.6 Estimation

Similar to DellaVigna et al. (2017), I fix $\eta = 1$ rather than being estimated; thus, the loss-aversion parameter λ can be interpreted as the overall weight on the losses. I also normalize the scale of the EV information acquisition cost shock $\sigma = 1$. The final set of parameters that I estimate is defined as $\theta := (\kappa, \beta, \lambda, \mu_{\xi}, \sigma_{\xi}, \nu)$ and summarized in Table 4.4.

TABLE 4.4: Description of the structural parameters

Parameter	Description
κ	Information acquisition cost
β	CRRA curvature
λ	Overall weight on loss (set $\eta = 1$)
μ_{ξ}	Mean of perception error
σ_{ξ}	Variance of perception error
ν	Effect of a reminder

¹³The intuition is similar to the Jensen inequality, I provide a simple illustration in Appendix C.4.

Estimation follows a simulated minimum-distance approach. Each bill j can be represented by tuple $(\tilde{G}_{j0}, \{c_{jt}, i_{jt}, r_{jt}\}_{t=1}^{T_j}, T_j, p_j)$, they are the bill's goal, daily energy consumption, daily information acquisition decision, whether the DM receives an Email reminder and the length of the bill and the per kWh price respectively.

To estimate θ , given a trial value of parameter vector $\tilde{\theta}$, I solve for the consumption and information acquisition policy function $c_t^*(\tilde{G}_{t-1}, \tilde{V}G_{t-1}), i_t^*(\tilde{G}_{t-1}, \tilde{V}G_{t-1})$ according to Section 4.4. Then for each bill j , starting with a goal G_j with a duration of T_j , I first draw S vectors unobservable ξ_j of length T_j according to $\tilde{\theta}$ and forward simulate a consumption and information acquisition path. Denote $c_{jts}(\tilde{\theta}), i_{jts}(\tilde{\theta})$ as the consumption and information acquisition for bill j , simulation s day t given the trial value of parameter vector $\tilde{\theta}$. I find the estimator by minimizing sum of squared distance between the model generated consumption and information acquisition decision $(c_{jts}(\tilde{\theta}), i_{jts}(\tilde{\theta}))$ and the observed consumption and information acquisition decision (c_{jt}, i_{jt}) for all observations. Specifically, the estimator is defined as:

$$\hat{\theta} = \underset{\theta}{\operatorname{argmin}} \sum_j \sum_t \left[\left(c_{jt} - \frac{1}{S} \sum_s c_{jts}(\theta) \right)^2 + \left(i_{jt} - \frac{1}{S} \sum_s i_{jts}(\theta) \right)^2 \right] \quad (4.29)$$

I provide a Monte Carlo simulation of the estimator in Appendix C.3, the simulation results show that the estimator is unbiased. Standard errors are calculated using bootstrapping, I used 800 (around 10% of the total sample) replications with replacement from a whole sample.

4.6.1 Other estimation details

I restrict the estimating sample to monthly bills (up to 45 days), as I follow [Nevo et al. \(2016\)](#) to ignore future discounting. I find it appropriate to do so only to monthly bills, where I may have to account for other behavioural biases such as present bias and projection bias for quarterly bills.

Since households receive weekly Emails, I set $q = 0.13$. Furthermore, households set weekly goals on the App, but they are billed monthly. When they interact with the App, they are informed about their estimated bill for the billing cycle instead of their (weekly) goal cycle. Furthermore, since households are billed monthly, it is natural to assume the gain-loss utility is realized at the end of the billing cycle rather than every week. Hence, I set T_j as the billing cycle length, and set $\tilde{G}_{0j}$ to tailor each billing cycle, that is, $\tilde{G}_{0j} = \text{goal}_{\text{week}}/7 \times T_j$, where $\text{goal}_{\text{week}}$ is the amount that the household enters into the

App. I am aware the literature on goal bracketing and its implications (Hsiaw, 2018), but I believe a billing cycle horizon is more realistic in this context.

Moreover, residential electricity is subject to non-linear pricing, where a typical pricing scheme includes a daily fixed and block pricing. However, it has been shown that consumers in this context are responsive to average prices instead of marginal prices (Ito, 2014), hence I construct average per kWh prices for each bill. Lastly, for technical reasons, there are some discrepancies between the bill's recorded usage (in kWh) and the usage from the smart meters. I discard bills that have a discrepancy that is more than 15 kWh in absolute terms. This leaves us with 8,553 bills from 531 households in the final sample, where I provide summary statistics in Table 4.5.

TABLE 4.5: Summary statistics for structural estimation sample

Variable	Mean	Std dev
Daily consumption (in kWh)	12.527	9.204
Bill total consumption (in kWh)	381.593	239.977
Bill in dollars	104.899	55.837
Total no. of information acquisition	3.309	3.886
Total no. of reminders	4.068	1.088
Billing cycle length	30.496	1.931
Prices (per kWh)	0.282	0.064
(Goal - Total consumption) in dollars	1.813	49.432
(Goal - Total consumption) if engaged	3.104	49.280
(Goal - Total consumption) if not engaged	-2.006	49.695
No. bills	8553	
No. households	531	

Notes: The level of observation in this summary statistics table is a bill. Here, I define *engaged* as acquiring information (checks an Email or logins the portal) at least once in the billing cycle.

4.7 Results

TABLE 4.6: Estimated structural parameters


Parameter	Description	Estimates	Std.Err
κ	Information acquisition cost	6.975	0.3775
β	CRRA curvature	0.275	0.0093
λ	Loss aversion	1.609	0.105
μ_{ξ}	Mean of belief error	-0.315	0.0104
σ_{ξ}	Variance of belief error	0.698	0.0724
ν	Effect of a reminder	3.938	0.1735
Fixed parameters			
η	Sensitivity to gain loss	1.000	-
σ	Scale of EV taste shock	1.000	-
	RMSE of c_{it}	8.035	
	RMSE of i_{it}	0.113	

Table 4.6 presents the structural parameter estimates. I only achieve a modest fit in my model for the following reasons. First, electricity consumption is highly heterogeneous across households and time. Since I only model heterogeneity in price and goals across households, it can be difficult to capture important features that affect electricity usage such as weather. Second, from Table 4.2, I find information acquisition may not necessarily depend on the states, hence information acquisition is likely to be driven by the EV shock in my model, resulting in a modest fit.

4.7.1 Parameters interpretation

To interpret and understand these behavioural parameters, I translate them into the effect on behaviours. κ represents the information acquisition cost that prevents households from acquiring information. Under a frictionless environment where κ is zero, holding all other parameters constant, an average household has a 54% probability to acquire information on a given day whereas the probability is 1.13% under the estimated parameter. Hence, the probability to acquire information reduces by 53% compared to a frictionless environment. Surprisingly, reducing information acquisition cost does not impact consumption on average, the reason is that μ_{ξ} is estimated to be small.

FIGURE 4.3: The effect of information acquisition costs




Notes: These figure presents the effect of information acquisition cost. I compare the probability of acquiring information and daily energy consumption from the estimated coefficients and $\kappa = 0$ (frictionless environment). The right panel is the effect of information acquisition cost on consumption c_{it} . The left panel is the effect of information acquisition cost on the probability of acquiring information i_{it}

μ_{ξ} presents the mean of the of the belief error that the household experiences if he does not acquire information on a day. -0.315 means on average, household underestimate their daily consumption by 0.315kWh , equivalent to approximately $\$0.1$ per day (assuming $\$0.3/\text{kWh}$). This accumulates to $\$2.835$ a month if the household never acquires information within the month. This only accounts for 2.7% of the total bill, which is relatively small in magnitude. This suggests that households have unbiased beliefs. While information acquisition corrects households' incorrect beliefs, their beliefs are not too far off from reality. Hence, I would not observe consumption changes after information acquisition in this subsample.

ν represents the reduction in information acquisition cost when the household receives a reminder, where $\nu = 0$ corresponds to the special case where a reminder has no effect on information acquisition cost when the household receives a reminder. When ν is zero, holding all other parameters constant, an average household has a 0.46% probability to acquire information on a given day whereas the probability is 1.19% under the estimated parameter. More importantly, the probability to acquire information when being reminded decreased from 1.19% to 0.46% when ν is reduced to zero.


FIGURE 4.4: The effect of a reminder



Notes: These figure presents the effect of a reminder. I compare the probability of acquiring information and daily energy consumption from the estimated coefficients and $\nu = 0$ (ineffective reminder). The right panel is the effect of Email reminders on consumption c_{it} . The left panel is the effect of Email reminders on the probability of acquiring information i_{it}

While information acquisition does not impact consumption, loss aversion does. $\lambda = 1$ corresponds to the special case where households are not loss averse, that is, the violation of the goal does not incur additional utility losses. Households on average, reduce their consumption by 13% compared to households that are not loss averse, holding all other factors constant.

FIGURE 4.5: The effect of loss aversion




Notes: These figure presents the effect of loss aversion. I compare the probability of acquiring information and daily energy consumption from the estimated coefficients and $\lambda = 0$ (consumers are not loss averse). The right panel is the effect of loss aversion on consumption c_{it} . The left panel is the effect of loss aversion on the probability of acquiring information i_{it}

4.7.2 Policy counterfactuals

I consider two types of policy counterfactuals: (i) giving households a 10% harder goal and (ii) sending reminder Emails to households every day. A tighter goal reduces energy consumption by approximately 3.6%, while sending reminders does not change energy usage. Households are more likely to acquire information when they receive reminders every day, whereas it does not affect information acquisition for a tighter goal.

FIGURE 4.6: Policy counterfactuals



Notes: The right panel is the effect of different policies on consumption c_{it} . The left panel is the effect of different policies on the probability of acquiring information i_{it}

I believe the tighter goal to be the more preferred policy to conserve energy in this context. As nudges are becoming a popular intervention to alter behaviour, policymakers and researchers are becoming more aware of the welfare cost associated with them. For example, [Allcott and Kessler \(2019\)](#) find the Home Energy Reports, a nudge that utilizes social comparison, incurs “non-energy costs” on the nudgee in an effort to reduce energy consumption. Furthermore, [Damgaard and Gravert, 2018](#)) finds charitable giving reminders give rise to non-negligible nuisance costs on the nudgee. Hence, since more reminders do not affect energy usage, the social cost that it incurs may lead to a welfare loss.

4.8 Conclusion

In this paper, I proposed a novel framework to study reference-dependence utility with costly information acquisition. I applied this model to the residential electricity industry.

I find that energy usage goals are effective to achieve energy conservation in the residential electricity context. I find that loss aversion is the main driver of the reduction in usage. Information acquisition cost plays a significant role in preventing households from acquiring information, and reminders reduce this cost. To my surprise, reducing information acquisition cost does not encourage further energy conservation. That is because households that set goals can track their progress quite unbiasedly in the absence of immediate feedback, so acquiring immediate feedback only reduces their uncertainty in beliefs, but it does not correct them.

Does that mean reminders and acquiring information is socially wasteful, or potentially detrimental in the residential energy sector? While information acquisition and reminders do not affect energy usage, I want to highlight a few caveats of my study.

First, information acquisition in the model has two functions, giving feedback on whether the household is on track to meet its goal and reducing uncertainty. [Allcott and Rogers \(2014\)](#) and Chapter 1 show bill arrivals decrease households' consumption as it acts as a cue that raises the salience of energy usage. My model does not capture such a channel. Hence, it is possible that reminders and information acquisition still impacts energy consumption and achieve conservation through increased salience.

Second, my model takes the reference point as exogenous. However, goal-setting is endogenous in this context: a household chooses whether to set a goal and the amount of the goal. I find goal setting households to be significantly different from the population. This is also supported by [Harding and Hsiaw \(2014\)](#) who finds that households who choose to set goals in the residential electricity context are generally more likely to express environmental interest and to be engaged in environmentally friendly actions. As I have discussed in Chapter 3, they may use the goal as a commitment device to conserve energy. So, it is important to emphasise that my empirical estimates are specific to this particular setting and there is no reason to think that the conservation estimates are representative of other populations, other institutional environments, or even other utility markets. However, at a broad level, my findings illustrate the relationship between goals, information, and biased beliefs, and I offer a cohesive framework to understand and analyse the intertwined relationship between these variables.

To overcome these limitations and to further understand goals, information, beliefs, and consumption in a more general setting, I direct future research to focus on the following. First, it is essential to understanding how goals are formed. Due to a lack of variation within a household in their goals, I cannot infer the goal-setting formation process. This is important to understand because the goal, as the reference point, is taken exogenously in the model. However, as previous studies suggest, goal-setting is highly likely to be endogenous to consumption, it has been shown to be a source of internal motivation to

address self-control problems ([Hsiaw, 2013](#)). In the broader literature of reference dependence utility, there are also different theories of reference point formation, and empirical evidence suggests that this process is highly heterogeneous ([Baillon et al., 2020](#)). For example, [Kőszegi and Rabin \(2006\)](#)'s preferred personal equilibrium characterizes the endogenous reference point with the decision maker's rational expectation. Hence, without a deeper understanding of how the goal is formed, the counterfactual where I exogenously change households' goals is less convincing, and I am reluctant to give policy recommendations to exogenously impose these reference points for households. Second, since the sample used to estimate my structural model is highly selective, I speculate that the unbiased belief results are likely to be specific to these energy-aware households. Hence, an RCT where the experimenter exogenously varies the goal for each household can test whether the results hold in a more general population.

Chapter 5

Conclusion

Behavioural interventions are becoming increasingly popular due to their cost-effectiveness. Contrary to conventional interventions such as taxes and price changes, these interventions do not significantly change individuals' economic incentives nor restrict individuals' choices ([Thaler and Sunstein, 2009](#)). We observe them being utilized in the residential electricity industry ([Abrahamse et al., 2005](#)). To implement effective policies, we must first have a thorough understanding of the consumers.

This thesis studies consumers' behaviour and behavioural interventions in the residential electricity industry. We see consumers exhibit behaviours that are consistent with "non-standard" preferences and beliefs in this industry¹. It gives us a better understanding of the consumers in this industry, so policymakers can design appropriate interventions to improve social welfare. On a broader level, it also makes several important contributions toward our understanding of inattention, information feedback, and goal-setting.

From chapter two, I examine whether individuals exhibit dynamic attention to their energy usage in the presence of intermittent billing. I ran a survey where I ask consumers to predict their forthcoming bill. I find consumers to be attentive to their energy usage for the first twelve days of the billing cycle. This finding is consistent with previous studies, that receiving the previous cycle's bill makes electricity usage salient early in the billing cycle. Since my test is expectation based, I find that not being attentive for the whole billing cycle results in bill shocks, where the consumer's expected payment does not align with the actual payment. To avoid bill shocks, policymakers can provide notifications and reminders in the middle of the billing cycle so consumers can be better prepared when they receive their bills.

¹Since the essays do not compare standard models against my proposed behavioural model, it is possible that alternate standard theories could also explain my empirical patterns. Ruling this out would require me to demonstrate my empirical patterns are inconsistent with other standard models.

The third and fourth chapters contribute to our understanding of how goal-setting with real-time information can affect individuals' behaviour. Traditionally, voluntary participation is the most common implementation of behavioural interventions, as nudges prohibit the restriction of individuals' choices. Hence, specific to my context, it is difficult to attribute changes in behaviour to goal-setting with only observational data since I expect individuals who self-select into goal-setting to be different from non-goal-setters. In chapter three, I estimate the effect of goal-setting in conjunction with real-time information feedback in the context of the residential electricity industry. I implemented a field experiment and randomly allocated consumers with an App that provides them with real-time feedback and an option to set a goal. To address selection into goal-setting, I employ a recently developed machine-learning causal inference method. The method semi-parametrically controls for household-time varying confounders. I find that (i) real-time information alone does not affect energy usage and (ii) goal-setting with real-time information feedback generates on average \$12 saving annually for a consumer. Without carefully controlling for household-time varying confounders would overstate the effect of goal-setting with real-time information by 64%. To promote energy conservation, policymakers can encourage households to set an energy saving target.

Chapter four aims to understand the mechanism that drives the estimated effect from chapter three. I propose and estimate a model of decision making, where the decision-maker exhibits reference-dependent utility and costly information acquisition. The decision-maker chooses daily consumption and information acquisition within a finite horizon, subject to a reference point. This novel framework allows us to understand the intertwined relationship between reference points, information friction, and decisions. It can isolate the effect of each feature on consumption and information acquisition. I apply this framework to the residential electricity industry, where the goal being the reference point of interest. Using the data from the experiment, I find that consumers' loss aversion accounts for 13% reduction in energy usage. While information acquisition cost plays a big role in reducing consumers' interaction with the App, not being fully informed in real-time does not affect energy usage. This is because consumers exhibit unbiased belief of their energy usage even if they do not acquire information, and they are already choosing consumption optimally based on their beliefs. Hence, since their beliefs are close to actuality, acquiring information does not 'correct beliefs' and make them re-optimize. However, I speculate this is due to sample selection, as goal-setters in my sample could be energy conscious consumers. In my counterfactual analysis, providing tighter goals reduces energy usage, while providing more Email reminders does not affect energy usage. Hence, since sending reminders to informed consumers does not affect consumption, it can potentially be socially wasteful.

Overall, I highlight the importance of inattention and non-standard preferences, especially in the residential electricity industry. If policymakers ignore their presence, this may lead to issues such as bill shocks. However, if policymakers understand and recognize the role that inattention and non-standard preferences play, they can use them in policy design within and outside the context of energy conservation.

Appendix A

Appendices for Chapter 2

A.1 Incentive compatibility from the pilot survey

In the pilot survey, the treatment group is incentivised to provide accurately beliefs. I implement a binary lottery procedure with quadratic scoring rule to achieve this. That is, the treatment group wins more lottery tickets to the grand prize, two iPad (RRP \$999 AUD each), iPhone 7 (RRP \$ 849). The control group is not incentivised to provide an accurate forecast: all have same probability to win the prize. I find no statistically difference between the treatment group and the control group.

TABLE A.1: Pilot survey results

	(1)	(2)	(3)	(4)
	Forecast error	Absolute Forecast error	Absolute Forecast error	Forecast error
Binary Lottery Procedure	-2.277 (3.824)	0.435 (3.227)	-0.586 (3.108)	-4.105 (3.738)
Constant	19.66*** (2.447)	40.83*** (2.049)	3.897 (9.515)	-4.004 (11.89)
Controls	No	No	Yes	Yes
Observations	1,295	1,295	1,295	1,295
R-squared	0.000	0.000	0.081	0.045

Robust standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

Notes: Treatment are household who are incentivised for accurate beliefs using the binary lottery procedure.

A.2 Survey email and questionnaire

SURVEY INVITATION EMAIL

Subject line: Make a guess about your electricity use to go in the draw

At _____ we're looking for new ways to help you better manage your electricity use. To do this, it would be useful to know a little more about your electricity use and lifestyle.

Please take a few moments to complete our short survey. As a thank you, you will be automatically entered into the draw for a chance to win one of two [apple product prize].

Once click on link

Wording for the link: "Complete Survey"

At the bottom of the email:

Terms and Conditions apply <insert link for T&Cs>. Entries close midnight <insert date>. Winners will be drawn on <insert date> at Level 33, _____ Melbourne and be published at _____ on <insert date>. Competition is only available to Victorian retail customers aged 18 years and over.

This email was sent to <insert email address> on behalf of _____ and The University of Melbourne. You received this email...<insert standard language here from Billcap>

FIGURE A.1: Survey Email

FIGURE A.2: Survey questionnaire

Survey

Thank you for taking the time to complete our short survey. Your answers will help ensure the information we share with you is useful and relevant.

1. Do you rent or own your home? Required
2. What type of home do you live in? Required
3. How many bedrooms are there in your home? Required
4. How many people live in your home? Required
5. How do you heat your home? Required
6. How do you heat your hot water? Required
7. How many air-conditioning units are in your home? Required
8. Do you have a swimming pool or spa? Required
9. Have you ever previously used a web-portal that showed you the electricity consumption recorded by your smart meter? Required
10. To the nearest dollar, what was your last monthly electricity bill? Required
11. To the nearest dollar, what do you expect your electricity bill to be next month? Required
12. To the nearest dollar, what do you expect your electricity bill to be in 2 months? Required
13. Have you had any unexpectedly high electricity bills in the past year? Required
14. If you answered yes to question 13, on average how much higher were the electricity bills than what you expected? Required
15. If you answered yes to question 13, can you recall which bills were unexpectedly high? Please check all that apply. Required

☐ 2015: Oct
☐ Nov
☐ Dec
☐ 2016: Jan
☐ Feb
☐ Mar
☐ Apr
☐ May
☐ Jun
☐ Jul
☐ Aug
☐ Does not apply

A.3 Balance test by survey waves

TABLE A.2: Balance test results - Wave 1

	10-days interval		Individual days	
	F-statistic	p-value	F-statistic	p-value
<i>(a) Survey data</i>				
Electricity and gas heating	0.053	0.949	0.801	0.774
Electricity heating	0.249	0.779	0.945	0.554
Gas heating	0.083	0.920	0.888	0.645
No of air conditioners	4.627	0.010	1.330	0.105
Uses central air conditioner	0.577	0.562	1.353	0.092
Apartment	0.729	0.482	1.025	0.428
House	1.709	0.181	0.918	0.597
Terrace	2.725	0.066	1.253	0.158
Townhouse	0.559	0.572	0.657	0.927
Owens the property	1.047	0.351	1.346	0.095
Rents the property	1.047	0.351	1.346	0.095
Has a pool or a spa	1.107	0.331	1.235	0.173
Solar hot water	1.116	0.328	0.463	0.995
Gas water heating	1.248	0.287	1.171	0.236
Electricity water heating	0.222	0.801	1.036	0.412
No of bedrooms	2.008	0.134	1.139	0.272
Household size	0.145	0.865	0.931	0.576
Experienced bill shock	0.052	0.949	0.900	0.626
<i>(b) Account data</i>				
Credit rating (out of 100)	3.075	0.046	1.084	0.343
Pays bill manually	1.926	0.146	0.774	0.810
<i>(c) Census data</i>				
Age	0.668	0.513	0.935	0.570
No of bedrooms	1.439	0.237	0.831	0.732
Prop of Australian citizens	0.873	0.418	0.615	0.953
Prop of houses	1.111	0.329	1.065	0.370
Median income	0.091	0.913	0.711	0.882
Prop of internet users	1.072	0.342	0.783	0.799
Median mortgage payment	0.217	0.805	0.970	0.514
Household size	0.489	0.613	0.906	0.617
Median rent payment	0.054	0.947	0.668	0.918
Prop of renters	1.499	0.223	0.734	0.857

Notes: This Table shows the results of our balance test. Column (2) - (3) represents the F-Statistics and p-value of the F-statistics for equation (2.2). Column (3) - (4) represents the F-Statistics and p-value of the F-statistics for equation (2.3). Results shows the day of the billing cycle when the consumer responds to our survey cannot predict the consumer's observable characteristics, implying they are comparable over the billing cycle. As results, this provide assurance to our identification strategies that consumers are exogenously exposed different amount of available information within the billing cycle.

TABLE A.3: Balance test results - Wave 2

	10-days interval		Individual days	
	F-statistic	p-value	F-statistic	p-value
<i>(a) Survey data</i>				
Electricity and gas heating	1.653	0.192	1.249	0.162
Electricity heating	1.197	0.302	0.981	0.497
Gas heating	0.337	0.714	0.822	0.745
No of air conditioners	0.852	0.427	1.312	0.116
Uses central air conditioner	0.317	0.728	1.152	0.258
Apartment	2.040	0.130	0.987	0.487
House	3.892	0.020	1.095	0.329
Terrace	0.572	0.564	1.039	0.407
Townhouse	2.201	0.111	1.108	0.311
Owns the property	0.232	0.793	1.011	0.450
Rents the property	0.232	0.793	1.011	0.450
Has a pool or a spa	2.623	0.073	0.862	0.686
Solar hot water	0.631	0.532	0.920	0.594
Gas water heating	0.227	0.797	1.542	0.028
Electricity water heating	0.157	0.854	1.092	0.333
No of bedrooms	2.239	0.107	0.596	0.963
Household size	0.265	0.767	0.908	0.614
Experienced bill shock	0.703	0.495	0.884	0.651
<i>(b) Account data</i>				
Credit rating (out of 100)	0.075	0.928	0.873	0.668
Pays bill manually	2.789	0.062	1.343	0.097
<i>(c) Census data</i>				
Age	0.591	0.554	0.615	0.953
No of bedrooms	1.089	0.337	0.844	0.712
Prop of Australian citizens	0.695	0.499	0.967	0.518
Prop of houses	0.679	0.507	1.118	0.299
Median income	1.014	0.363	1.457	0.049
Prop of internet users	0.053	0.949	0.896	0.633
Median mortgage payment	0.670	0.512	1.385	0.076
Household size	1	0.368	0.823	0.744
Median rent payment	1.176	0.309	0.577	0.971
Prop of renters	0.583	0.558	1.035	0.413

Notes: This Table shows the results of our balance test. Column (2) - (3) represents the F-Statistics and p-value of the F-statistics for equation (2.2). Column (3) - (4) represents the F-Statistics and p-value of the F-statistics for equation (2.3). Results shows the day of the billing cycle when the consumer responds to our survey cannot predict the consumer's observable characteristics, implying they are comparable over the billing cycle. As results, this provide assurance to our identification strategies that consumers are exogenously exposed different amount of available information within the billing cycle.

TABLE A.4: Balance test results - Wave 3

	10-days interval		Individual days	
	F-statistic	p-value	F-statistic	p-value
<i>(a) Survey data</i>				
Electricity and gas heating	1.653	0.192	1.203	0.207
Electricity heating	1.197	0.302	0.859	0.686
Gas heating	0.337	0.714	0.671	0.912
No of air conditioners	0.852	0.427	0.568	0.972
Uses central air conditioner	0.317	0.728	0.967	0.517
Apartment	2.040	0.130	0.948	0.547
House	3.892	0.020	1.172	0.238
Terrace	0.572	0.564	0.783	0.794
Townhouse	2.201	0.111	1.437	0.059
Owns the property	0.232	0.793	0.969	0.514
Rents the property	0.232	0.793	0.969	0.514
Has a pool or a spa	2.623	0.073	1.334	0.106
Solar hot water	0.631	0.532	0.921	0.590
Gas water heating	0.227	0.797	0.791	0.784
Electricity water heating	0.157	0.854	0.890	0.638
No of bedrooms	2.239	0.107	0.949	0.545
Household size	0.265	0.767	1.026	0.427
Experienced bill shock	0.703	0.495	1.145	0.270
<i>(b) Account data</i>				
Credit rating (out of 100)	0.919	0.399	0.680	0.905
Pays bill manually	0.358	0.699	1.072	0.361
<i>(c) Census data</i>				
Age	0.666	0.514	1.276	0.145
No of bedrooms	0.152	0.859	1.033	0.416
Prop of Australian citizens	0.722	0.486	0.900	0.624
Prop of houses	0.776	0.460	0.984	0.491
Median income	1.847	0.158	1.369	0.088
Prop of internet users	0.736	0.479	1.486	0.043
Median mortgage payment	1.483	0.227	1.219	0.192
Household size	0.074	0.929	1.019	0.438
Median rent payment	2.223	0.109	1.418	0.066
Prop of renters	0.718	0.488	1.300	0.127

Notes: This Table shows the results of our balance test. Column (2) - (3) represents the F-Statistics and p-value of the F-statistics for equation (2.2). Column (3) - (4) represents the F-Statistics and p-value of the F-statistics for equation (2.3). Results shows the day of the billing cycle when the consumer responds to our survey cannot predict the consumer's observable characteristics, implying they are comparable over the billing cycle. As results, this provide assurance to our identification strategies that consumers are exogenously exposed different amount of available information within the billing cycle.

A.4 Derivation and inferences of turning points

A.4.1 Quadratic specification

$RMSE_{is} = \beta_0 + \beta_1 Day_{is} + \beta_2 Day_{is}^2 + \dots$ To find the turning point, I take the derivative of $RMSE_{is}$ with respect to Day_{is} and set to 0:

$$\frac{\partial RMSE_{is}}{\partial Day_{is}} : \beta_1 + 2\beta_2 Day_{is} = 0$$

$$\hat{\theta}_{\text{Quadratic}} = -\frac{\hat{\beta}_1}{2\hat{\beta}_2}$$

The confidence intervals constructed for $\hat{\theta}$ using the Delta method is to estimate its variance by use of a first order Taylor series expansion to linearise the non-linear function. The $100 \times (1 - \alpha)\%$ confidence interval for $\hat{\theta}$ is defined by [Rao et al. \(1973\)](#):

$$\hat{\theta} \pm t_{\alpha/2} \sqrt{\frac{\hat{\sigma}_1^2 \hat{\beta}_2^2 - 2\hat{\beta}_1 \hat{\beta}_2 \hat{\sigma}_{12} + \hat{\beta}_1^2 \hat{\sigma}_2^2}{4\hat{\beta}_2^4}}$$

where $\hat{\sigma}_i^2$ is the estimated variance of $\hat{\beta}_i$, and $\hat{\sigma}_{ij}$ is the estimated covariance between $\hat{\beta}_i$ and $\hat{\beta}_j$

A.4.2 Cubic specification

$RMSE_{is} = \beta_0 + \beta_1 Day_{is} + \beta_2 Day_{is}^2 + \beta_3 Day_{is}^3 \dots$ To find the turning point, I take the derivative of $RMSE_{is}$ with respect to Day_{is} and set to 0:

$$\frac{\partial RMSE_{is}}{\partial Day_{is}} : \beta_1 + 2\beta_2 Day_{is} + 3\beta_3 Day_{is}^2 = 0$$

$$\hat{\theta}_{\text{Cubic}} = -\frac{2\beta_2 \pm \sqrt{4\beta_2^2 - 12\beta_1\beta_3}}{6\beta_3}$$

The confidence intervals constructed for $\hat{\theta}$ using the Delta method is to estimate its variance by use of a first order Taylor series expansion to linearise the non-linear function. The $100 \times (1 - \alpha)\%$ confidence interval for $\hat{\theta}$ is defined by [Rao et al. \(1973\)](#):

$$\hat{\theta}_i \pm t_{\alpha/2} d_1 \sigma_1^2 + 2d_1 d_2 \sigma_{12} + 2d_1 d_3 \sigma_{13} + d_2 \sigma_2^2 + 2d_2 d_3 \sigma_{23} + d_3 \sigma_3^2$$

where $\hat{\sigma}_i^2$ is the estimated variance of $\hat{\beta}_i$, and $\hat{\sigma}_{ij}$ is the estimated covariance between $\hat{\beta}_i$ and $\hat{\beta}_j$. $d_1 = \frac{\partial \theta_i}{\partial \beta_1}, d_2 = \frac{\partial \theta_i}{\partial \beta_2}, d_3 = \frac{\partial \theta_i}{\partial \beta_3}, i = 1, 2$

A.4.3 Quartic specification

There is no close-form solution for the turning point, hence I solve for the turning point numerically.

A.5 Robustness checks

TABLE A.5: Robustness checks for for Equation (2.1)

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	RMSE (% of total bill)	RMSE	RMSE	RMSE	RMSE	RMSE	RMSE	RMSE
Day	-0.014*** (0.005)	-3.425** (1.590)	-0.138 (0.084)	-0.148 (0.096)	-1.588*** (0.327)	-1.603*** (0.378)	-3.372*** (0.773)	-3.158*** (0.883)
Day ²	0.000*** (0.000)	0.222 (0.221)			0.049*** (0.010)	0.049*** (0.012)	0.201*** (0.060)	0.182*** (0.069)
Day ³		-0.005 (0.011)					-0.003*** (0.001)	-0.003** (0.002)
Day ⁴		0.000 (0.000)						
Constant	0.517*** (0.032)	55.049*** (3.363)	44.369*** (1.500)	-32.982*** (8.322)	51.263*** (2.265)	-25.206*** (8.614)	55.315*** (2.846)	-21.852** (8.675)
Observations	9,254	9,353	9,353	7,262	9,353	7,262	9,353	7,262
R-squared	0.017	0.007	0.000	0.071	0.003	0.073	0.004	0.074
Controls	N	N	N	Y	N	N	N	Y
Survey FE	Y	Y	N	N	N	N	N	N
Mean of dep. variable	0.428	0.428	42.35	43.64	42.35	43.64	42.35	43.64
1st Turning point	17.60	12.01			16.24	16.31	12.11	12.49
Standard error (first)	1.312				0.852	0.964	1.193	1.575
2nd Turning point		28.90					27.35	28.27
3rd Turning point		86.88						
Standard error (second)							3.068	4.623

Robust standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

A.6 Testing differences in turning points for two subgroups

A.6.1 Quadratic specification

We regress the following equation, we drop subscripts for the convenience of notation. Denote Z as the group indicator (analogous to *ManualPayment* in our specification).

$$RMSE = \beta_0 + \beta_1 Day + \beta_2 Day^2 + \beta_3 Day \times Z + \beta_4 Day^2 \times Z + \beta_5 Z + \varepsilon$$

Following Appendix A.4, turning point for $Z = 1$ is

$$\theta_{Z=1} = -\frac{\beta_3 + \beta_5}{2(\beta_1 + \beta_4)}$$

and turning point for $Z = 0$ is

$$\theta_{Z=0} = -\frac{\beta_1}{2\beta_3}$$

Hence, the following hypothesis test tests for the difference in the first turning point,

$$\begin{aligned} H_0 : -\frac{\beta_1 + \beta_5}{2(\beta_3 + \beta_4)} + \frac{\beta_1}{2\beta_3} &= 0 \\ H_1 : -\frac{\beta_1 + \beta_5}{2(\beta_3 + \beta_4)} + \frac{\beta_1}{2\beta_3} &\neq 0 \end{aligned}$$

I calculate the standard error using the Delta method.

A.6.2 Cubic specification

Following the same notation, I regress the following equation,

$$RMSE = \beta_0 + \beta_1 Day + \beta_2 Day^2 + \beta_3 Day^3 + \beta_4 Day \times Z + \beta_5 Day^2 \times Z + \beta_6 Day^3 \times Z + \beta_7 Z + \varepsilon$$

Following Appendix A.4, turning point for $Z = 1$ is

$$\theta_{Z=1} = \frac{-2(\beta_2 + \beta_5) \pm \sqrt{4(\beta_2 + \beta_5)^2 - 12(\beta_1 + \beta_4)(\beta_1 + \beta_4)}}{6(\beta_3 + \beta_6)}$$

and turning point for $Z = 0$ is

$$\theta_{Z=0} = \frac{-2\beta_2 \pm \sqrt{4\beta_2^2 - 12\beta_3\beta_1}}{6\beta_3}$$

Hence, the following hypothesis test tests for the difference in the first turning point:

$$\begin{aligned} H_0 : \frac{-2(\beta_2 + \beta_5) \pm \sqrt{4(\beta_2 + \beta_5)^2 - 12(\beta_1 + \beta_4)(\beta_1 + \beta_4)}}{6(\beta_3 + \beta_6)} - \frac{-2\beta_2 \pm \sqrt{4\beta_2^2 - 12\beta_3\beta_1}}{6\beta_3} &= 0 \\ H_1 : \frac{-2(\beta_2 + \beta_5) \pm \sqrt{4(\beta_2 + \beta_5)^2 - 12(\beta_1 + \beta_4)(\beta_1 + \beta_4)}}{6(\beta_3 + \beta_6)} - \frac{-2\beta_2 \pm \sqrt{4\beta_2^2 - 12\beta_3\beta_1}}{6\beta_3} &\neq 0 \end{aligned}$$

I calculate the standard error using the Delta method.

Appendix B

Appendices for Chapter 3

B.1 App usage overtime

To investigate heterogeneity in App engagement over time non-parametrically, I estimate a *Group based trajectory model*. It is a finite mixture model for panel data that is widely adopted in other social and clinical science research (see [Nagin and Odgers, 2010](#), for a review) to investigate different (unobservable) subgroup in treatment adherence over time. The conceptual aim of this analysis is to identify clusters of individuals with similar usage patterns over time. This method is attractive because I can identify heterogeneity (in engage patterns over time) semi-parametrically, that is, without inclusion of any observable characteristics as predictors.

The model can be summarised in the following equations:

$$Engage_{it} = \mathbf{1}[\beta_{j,0} + \sum_{k=1}^K \beta_{j,k} Time_{it}^k + \varepsilon_{it} > 0] \quad (B.1)$$

$$\varepsilon_{it} \sim \text{EV1} \quad (B.2)$$

$$p^j(Engage_{it}) = \left(\frac{\exp(\beta_{j,0} + \sum_{k=1}^{K_j} \beta_{j,k} Time_{it}^k)}{1 + \exp(\beta_{j,0} + \sum_{k=1}^{K_j} \beta_{j,k} Time_{it}^k)} \right)^{Engage_{it}} \left(\frac{1}{1 + \exp(\beta_{j,0} + \sum_{k=1}^{K_j} \beta_{j,k} Time_{it}^k)} \right)^{1-Engage_{it}} \quad (B.3)$$

$$P^j(Engage_i) = \left[\prod_{t=1}^{\tau_i-1} p^j(Engage_{it} | w_{it} = 1)(1 - \theta_t^j) \right] \theta_{\tau_i}^j \quad (B.4)$$

$$\theta_t^j = \frac{1}{1 + \exp(\gamma_0 + \gamma_j Season_{it})} \quad (B.5)$$


$$P(Engage_i) = \sum_{j=1}^J \pi_j P^j(Engage_i) \quad (B.6)$$

$$\sum_{j=1}^J \pi_j = 1 \quad (B.7)$$

$$L(\boldsymbol{\pi}, \boldsymbol{\beta}) = \prod_{i=1}^N P(Engage_i) \quad (B.8)$$

where $Engage_{it}$ is a binary indicator that denotes whether household i engaged with the platform in month t . I define technology engagement as a binary indicator that takes on the value 1 if the household either logs in the portal or checks an Email. We are interested in $\boldsymbol{\pi} = \{\pi_1, \pi_2, \dots, \pi_J\}$, and $\boldsymbol{\beta} = \{\beta_{1,0}, \beta_{1,1}, \dots, \beta_{J,K_j}\}$. π_j can be interpreted as the probability of a individual being in latent class j . $\beta_j = \{\beta_{j,1} \dots \beta_{j,K_j}\}$ is a polynomial approximation of the usage pattern of the latent class j .

FIGURE B.1: App usage overtime




Note: This figure plots the predicted usage patterns for our preferred specification. Y-axis represents the predicted probability of using the App, X-axis represents number of months since treatment.

I plot the predicted usage patterns for our preferred specification in Figure B.1. There are five subgroups in our samples: always users (30%), medium users (22%), never users (19%), decayers (15%) and learners (14%). Always users has a high probability (more than 90%) to engages with the platform every month, independent of how long he has been using the platform. Medium users has a medium probability (30% – 50%) , to engages with the platform every month, independent of how long he has been using the platform. Never users has a low probability (less than 10%) to engages with the platform every month, independent of how long he has been using the platform. Decayer’s probability to engages with the platform falls over time, from 90% in the first month to 50% after two years. Learner’s probability to engages with the platform rise over time, from 40% in the first month to 80% after two years.

B.2 Geographic locations of our sample and weather stations

FIGURE B.2: Spatial distribution of households in the sample and weather stations



Note: This figure shows the spatial distribution of households in our sample and weather stations. Shaded areas are SA1 with at least one household in the sample, red crosses are the weather stations

B.3 Regression results for diff-in-diff models by 3 hours block

TABLE B.1: Regression results for diff-in-diff models by 3 hours block

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	0:00 - 3:00	3:00 - 6:00	6:00 - 9:00	9:00 - 12:00	12:00 - 15:00	15:00 - 18:00	18:00 - 21:00	21:00 - 24:00
App	-0.001 (0.004)	-0.001 (0.003)	-0.002 (0.004)	-0.005 (0.004)	-0.007* (0.004)	-0.009** (0.004)	-0.009** (0.005)	-0.001 (0.004)
App-and-Goal	-0.048** (0.020)	-0.033** (0.017)	-0.007 (0.024)	-0.023 (0.022)	-0.023 (0.021)	-0.046* (0.025)	-0.042 (0.028)	-0.058*** (0.021)
Observations	57,303,496	57,303,496	57,303,496	57,303,496	57,303,496	57,303,496	57,303,496	57,303,496
R-squared	0.480	0.409	0.470	0.431	0.404	0.425	0.460	0.459

Robust standard errors in parentheses

*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

B.4 Regression results for ML diff-in-diff models by 3 hours block

TABLE B.2: Regression results for ML diff-in-diff models by 3 hours block

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	0:00 - 3:00	3:00 - 6:00	6:00 - 9:00	9:00 - 12:00	12:00 - 15:00	15:00 - 18:00	18:00 - 21:00	21:00 - 24:00
App	-0.001 (0.001)	0.001 (0.001)	0.001 (0.002)	0.000 (0.002)	-0.000 (0.002)	-0.002 (0.002)	-0.002 (0.002)	0.001 (0.002)
App-and-Goal	-0.010 (0.007)	-0.015*** (0.006)	-0.017** (0.008)	-0.000 (0.008)	0.008 (0.009)	-0.000 (0.011)	-0.018 (0.012)	-0.018* (0.009)
posttrain	0.015*** (0.002)	0.011*** (0.002)	0.041*** (0.003)	0.029*** (0.003)	0.027*** (0.003)	0.033*** (0.004)	0.074*** (0.004)	0.039*** (0.003)
Observations	49,256,007	49,255,511	49,251,061	49,264,127	49,281,058	49,278,135	49,261,524	49,256,551
R-squared	0.083	0.093	0.079	0.055	0.053	0.062	0.080	0.081

Robust standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

B.5 Robustness checks for ML diff-in-diff models by 3 hours block

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	0:00 - 3:00	3:00 - 6:00	6:00 - 9:00	9:00 - 12:00	12:00 - 15:00	15:00 - 18:00	18:00 - 21:00	21:00 - 24:00
App	-0.001 (0.001)	0.001 (0.001)	0.001 (0.002)	0.001 (0.002)	0.000 (0.002)	-0.002 (0.002)	-0.002 (0.002)	0.001 (0.002)
App-and-Goal	-0.011* (0.006)	-0.016*** (0.006)	-0.016** (0.008)	-0.001 (0.008)	0.008 (0.009)	-0.002 (0.011)	-0.017 (0.012)	-0.017* (0.009)
posttrain	0.015*** (0.002)	0.011*** (0.002)	0.040*** (0.003)	0.030*** (0.003)	0.029*** (0.003)	0.035*** (0.004)	0.076*** (0.004)	0.040*** (0.003)
Observations	49,840,319	49,839,757	49,835,367	49,848,979	49,866,011	49,862,753	49,846,124	49,841,163
R-squared	0.084	0.094	0.080	0.055	0.054	0.063	0.081	0.081

Robust standard errors in parentheses

*** p<0.01, ** p<0.05, * p<0.1

Appendix C

Appendices for Chapter 4

C.1 Variables descriptions for the model

TABLE C.1: Variables description

Variable	Description
Initial conditions (given)	
$\tilde{G}_0$	Goal for the billing cycle
$\tilde{V}G_0$	variance on the belief of goal for the billing cycle, equals 0
Deterministic state variables	
G_{t-1}	belief of the <i>remainder of the goal</i> before he violate this goal at time t . Normally distributed with mean $\tilde{G}_{t-1}$, variance $\tilde{V}G_{t-1}$
$\tilde{G}_{t-1}$	mean of the belief at time t
$\tilde{V}G_{t-1}$	variance of the belief at time t
Stochastic state variables	
r_t	equals 1 if the DM receives an Email reminder on day t , realization of Bernoulli(q)
ε_t	information acquisition cost shock, realization of $EV1(\sigma)$
$\tilde{\zeta}_t$	belief noise, realization of $N(\mu_{\tilde{\zeta}}, \sigma_{\tilde{\zeta}})$
Choice variables	
c_t	electricity usage on day t (continuous)
i_t	information acquisition decision (binary)

Functions

$GL(\cdot)$	gain-loss utility
$EV1(\sigma)$	Type-1 extreme value distribution with scaling parameter σ
$N(\mu, \sigma)$	Normal distribution with mean μ , variance σ^2

Behavioural parameters

q	probability of receiving an Email reminder
κ	information acquisition cost
β	CRRA curvature
λ	loss aversion
η	sensitivity to gain-loss utility function
μ_{ξ}	mean of perception error
σ_{ξ}	variance of perception error
ν	effect of a reminder

C.2 Modelling present-bias

Early work by [Rust \(1987\)](#) shows that the discount factor is usually not identified in dynamic discrete choice models. There is also a growing literature to understand the identification of present biased discount functions (e.g. [O'Donoghue and Rabin \(1999\)](#)) in Dynamic Discrete Choice Models ([Abbring et al., 2018](#)). One could modify the consumption flow utility Equation (4.3) to

$$u(c_t) = \frac{c_t^{1-\beta} - 1}{1-\beta} \quad (\text{C.1})$$

and introduce another state variable that represents total bill up to date t

$$\bar{C}_t = \bar{C}_{t-1} + pc_t \quad (\text{C.2})$$

and modify the terminal period utility from Equation (4.7) to

$$GL(\tilde{G}_T, \tilde{V}_{G_T}, \bar{C}_T) = \begin{cases} \eta G_T + \bar{C}_T & \text{if } G_T > 0 \\ \eta \lambda G_T + \bar{C}_T & \text{otherwise} \end{cases} \quad (\text{C.3})$$

$$\text{where } G_T \sim N(\tilde{G}_T, \tilde{V}_{G_T}) \quad (\text{C.4})$$

with this setup, the DM's flow utility only consist of the enjoyment of the electricity usage, but incur disutility of the bill at the terminal period. We believe this setup lacks the exclusion variables that affect the states transition over time, but do not affect the flow utility function, which is essential to identify the discounting function. As we observe c_t enters both (C.1) and (C.2).

C.3 Monte Carlo study

In this section, we perform a Monte Carlo study to understand the statistical properties of our estimator, and check whether we are able to recover the structural parameters if they were known. In each Monte Carlo run, we simulate 1,500 artificial bills using parameters from Table C.2, setting from Table C.3 and estimate the structural parameters following Section 4.6. We repeat this exercise 100 times and report the sampling distribution of the parameters in Figure C.1 and Table C.4.

TABLE C.2: Simulation baseline parameters

Parameter	Description	Value
κ	Information acquisition cost	2
β	CRRA curvature	1
λ	Loss aversion	5
η	Sensitivity to gain loss	1
μ_{ξ}	Mean of perception error	0
σ_{ξ}	Variance of perception error	2
σ	Scale of EV taste shock	1
ν	Effect of the nudge	1

TABLE C.3: Simulation settings

Notation	Description	Value
T	Number of periods	10
-	# grid points over assets	50
A_1	Minimum level - belief of reminder mean state	-80
A_f	Maximum level - belief of reminder mean state	50
-	Maximum level of belief variance state	30
-	# grids for belief variance state	30
-	# quadrature points used in calculation of expectations	50
S	# simulations draw to integrate ξ_t	200
p	price	0
N	# of bills	1500
p_n	Probability of being nudged	0

FIGURE C.1: Sampling distribution for Monte Carlo study




TABLE C.4: Comparing Monte Carlo results with actual parameters

Parameter	Actual	Mean of estimated
κ	2.000	2.029
β	0.800	0.795
λ	5.000	5.225
η	1.000	1.000
μ_{ζ}	0.000	0.103
σ_{ζ}	2.000	2.008
σ	1.000	1.000
ν	1.000	1.074

From Figure C.1 and Table C.4, our estimator is unbiased, and on average the estimated value is very close to the actual parameter.


C.4 Loss aversion and information acquisition

For this illustration, suppose the DM is loss averse ($\lambda > 1$). His gain loss function ($GL(\cdot)$) is depicted by the bold line in Figure C.2.

His belief G that takes on -10 and 10 with equal probability of 0.5. We denote that as point A and B respective in Figure C.2. We can think of this as a gamble. The expected value

of this gamble is $0.5 \times -10 + 0.5 \times 10 = 0$. That means, on average, he achieves his goal. However, as illustrated in Figure C.2, the expected gain loss of this gamble $EGL(G)$ is negative. This is because the DM is loss aversion, hence the gain-loss function is concave. We can also think about the information acquisition cost κ as the certainty equivalence, denoted by $CE(G)$ in Figure C.2.

FIGURE C.2: Illustration of the gain-loss utility



Bibliography

- Abadie, A. and Kasy, M. (2019). Choosing among regularized estimators in empirical economics: The risk of machine learning. *Review of Economics and Statistics*, 101(5):743–762.
- Abbring, J. H., Daljord, Ø., and Iskhakov, F. (2018). Identifying present-biased discount functions in dynamic discrete choice models. Technical report, Working paper, Tilburg University.[472].
- Abeler, J., Falk, A., Goette, L., and Huffman, D. (2011). Reference points and effort provision. *American Economic Review*, 101(2):470–92.
- Abrahamse, W., Steg, L., Vlek, C., and Rothengatter, T. (2005). A review of intervention studies aimed at household energy conservation. *Journal of Environmental Psychology*, 25(3):273–291.
- Agarwal, S., Goette, L., Sing, T. F., Staake, T., and Tiefenbeck, V. (2017). The role of goals and real-time feedback in resource conservation: Evidence from a large-scale field experiment. *National University of Singapore, Singapore*.
- AGL Energy (2020). AGL energy (version 6.14.1)[mobile application software]. Retrieved from https://play.google.com/store/apps/details?id=agl.digital.mobile&hl=en_AU&gl=US.
- Allcott, H. and Kessler, J. B. (2019). The welfare effects of nudges: A case study of energy use social comparisons. *American Economic Journal: Applied Economics*, 11(1):236–76.
- Allcott, H. and Rogers, T. (2014). The short-run and long-run effects of behavioral interventions: Experimental evidence from energy conservation. *American Economic Review*, 104(10):3003–37.
- Allcott, H. and Taubinsky, D. (2015). Evaluating behaviorally motivated policy: Experimental evidence from the lightbulb market. *American Economic Review*, 105(8):2501–38.
- Allen, E. J., Dechow, P. M., Pope, D. G., and Wu, G. (2016). Reference-dependent preferences: Evidence from marathon runners. *Management Science*, 63(6):1657–1672.

- Apple (2018). Close your rings. Retrieved from <https://www.apple.com/au/watch/close-your-rings/>.
- Attari, S. Z., Gowrisankaran, G., Simpson, T., and Marx, S. M. (2014). Does information feedback from in-home devices reduce electricity use? evidence from a field experiment. Technical report, National Bureau of Economic Research.
- Australian Bureau of Statistics (2013). Household Energy Consumption Survey, Australia: Summary of Results, 2012.
- Australian Energy Market Commission (2007). Rule change proposal (jurisdictional derogation) advanced metering infrastructure. *Department of Primary Industries, Victorian Government*.
- Australian Energy Regulator (2016). AER releases findings on smart meter spending. Retrieve from <https://www.aer.gov.au/news-release/aer-releases-findings-on-smart-meter-spending>.
- Australian Energy Regulator (2017). Electricity and gas bill benchmarks for residential customers 2017.
- Baillon, A., Bleichrodt, H., and Spinu, V. (2020). Searching for the reference point. *Management Science*, 66(1):93–112.
- Bainbridge, A. (2014). Survey finds 1 in 8 Australians cannot afford to pay electricity bill. *ABC News*. <http://www.abc.net.au/news/2014-10-13/1-in-8-australians-cant-pay-electricity-bill/5805886>.
- Becker, L. J. (1978). Joint effect of feedback and goal setting on performance: A field study of residential energy conservation. *Journal of Applied Psychology*, 63(4):428.
- Burlig, F., Knittel, C., Rapson, D., Reguant, M., and Wolfram, C. (2020). Machine learning from schools about energy efficiency. *Journal of the Association of Environmental and Resource Economists*, 7(6):1181–1217.
- Byrne, D. P., Nauze, A. L., and Martin, L. A. (2018). Tell me something I don't already know: Informedness and the impact of information programs. *Review of Economics and Statistics*, 100(3):510–527.
- Calzolari, G. and Nardotto, M. (2017). Effective reminders. *Management Science*, 63(9):2915–2932.
- Carbon Credits (Carbon Farming Initiative-Aggregated Small Energy Users) Methodology Determination 2015. (Cth).

Chan, E. Are consumers benefiting from smart meters? evaluating the long run utilization of an electricity feedback platform. *Working paper*.

Chetty, R., Looney, A., and Kroft, K. (2009). Salience and taxation: Theory and evidence. *American Economic Review*, 99(4):1145–77.

CITA (2011). CTIA, Federal Communications Commission and Consumers Union Announce Free Alerts to Help Consumers Avoid Unexpected Overage Charges. <https://www.ctia.org/industry-data/press-releases-details/press-releases/ctia-federal-communications-commission-and-consumers-union-announce-free-alerts-to-help-consumers-avoid-unexpected-overage-charges>

Clark, D., Gill, D., Prowse, V., and Rush, M. (2020). Using goals to motivate college students: Theory and evidence from field experiments. *Review of Economics and Statistics*, pages 1–45.

Clean Energy Regulator (2020). Emission reduction fund project register. Retrieve from <http://www.cleanenergyregulator.gov.au/ERF/project-and-contracts-registers/project-register>.

Crawford, V. P. and Meng, J. (2011). New york city cab drivers' labor supply revisited: Reference-dependent preferences with rational-expectations targets for hours and income. *American Economic Review*, 101(5):1912–32.

Damgaard, M. T. and Gravert, C. (2018). The hidden costs of nudging: Experimental evidence from reminders in fundraising. *Journal of Public Economics*, 157:15–26.

DellaVigna, S., Lindner, A., Reizer, B., and Schmieder, J. F. (2017). Reference-dependent job search: Evidence from hungary. *The Quarterly Journal of Economics*, 132(4):1969–2018.

Farber, H. S. (2008). Reference-dependent preferences and labor supply: The case of new york city taxi drivers. *American Economic Review*, 98(3):1069–82.

Faruqui, A., Sergici, S., and Sharif, A. (2010). The impact of informational feedback on energy consumption—a survey of the experimental evidence. *Energy*, 35(4):1598–1608.

FitNow Inc. (2020). Ocalorie counter by lose it! for diet & weight loss (version 12.7.101)[mobile application software]. Retrieved from <https://play.google.com/store/apps/details/?id=com.fitnow.loseit>.

Fowlie, M., Wolfram, C., Spurlock, C. A., Todd, A., Baylis, P., and Cappers, P. (2017). Default effects and follow-on behavior: evidence from an electricity pricing program. Technical report, National Bureau of Economic Research.

- Gilbert, B. and Zivin, J. G. (2014). Dynamic salience with intermittent billing: Evidence from smart electricity meters. *Journal of Economic Behavior & Organization*, 107:176–190.
- Goerg, S. J. and Kube, S. (2012). Goals (th) at work—goals, monetary incentives, and workers’s performance. *MPI Collective Goods Preprint*, (2012/19).
- Grubb, M. D. (2014). Consumer inattention and bill-shock regulation. *The Review of Economic Studies*, 82(1):219–257.
- Grubb, M. D. and Osborne, M. (2014). Biased beliefs, learning, and bill shock. *The American Economic Review*, 105(1):234–271.
- Guthridge, G., Burns, A., and Pelotti, P. (2012). Actionable insights for the new energy consumer.
- Harding, M. and Hsiaw, A. (2014). Goal setting and energy conservation. *Journal of Economic Behavior & Organization*, 107:209–227.
- Harrison, G. W., Martínez-Correa, J., and Swarthout, J. T. (2014). Eliciting subjective probabilities with binary lotteries. *Journal of Economic Behavior & Organization*, 101:128–140.
- Heath, C., Larrick, R. P., and Wu, G. (1999). Goals as reference points. *Cognitive Psychology*, 38(1):79–109.
- Hortaçsu, A., Madanizadeh, S. A., and Puller, S. L. (2017). Power to choose? an analysis of consumer inertia in the residential electricity market. *American Economic Journal: Economic Policy*, 9(4):192–226.
- Hsiaw, A. (2013). Goal-setting and self-control. *Journal of Economic Theory*, 148(2):601–626.
- Hsiaw, A. (2018). Goal bracketing and self-control. *Games and Economic Behavior*, 111:100–121.
- ING (2018). Savings goal calculator. Retrieved from <https://www.ing.com.au/savings/calculators/savings-goal.html>.
- Iskhakov, F., Jørgensen, T. H., Rust, J., and Schjerning, B. (2017). The endogenous grid method for discrete-continuous dynamic choice models with (or without) taste shocks. *Quantitative Economics*, 8(2):317–365.
- Ito, K. (2014). Do consumers respond to marginal or average price? evidence from non-linear electricity pricing. *American Economic Review*, 104(2):537–63.
- Jessoe, K. and Rapson, D. (2014). Knowledge is (less) power: Experimental evidence from residential energy use. *American Economic Review*, 104(4):1417–38.

- Karlan, D., McConnell, M., Mullainathan, S., and Zinman, J. (2016). Getting to the top of mind: How reminders increase saving. *Management Science*, 62(12):3393–3411.
- Keefer, Q. and Rustamov, G. (2015). Limited attention in residential energy markets: a regression discontinuity approach. *Empirical Economics*, pages 1–25.
- Kőszegi, B. and Rabin, M. (2006). A model of reference-dependent preferences. *The Quarterly Journal of Economics*, 121(4):1133–1165.
- Laskie, A. (2019). Melbourne heatwave: Eleven ways to keep your cool. *The Age*. <https://www.theage.com.au/national/victoria/melbourne-heatwave-eleven-ways-to-keep-your-cool-20191229-p53nh0.html>.
- McFadden, D. (1973). Conditional logit analysis of qualitative choice behavior.
- Mey, A. and Hoff, S. (2017). U.S. Energy Information Administration - EIA - Independent Statistics and Analysis.
- Nagin, D. S. and Odgers, C. L. (2010). Group-based trajectory modeling in clinical research. *Annual Review of Clinical Psychology*, 6:109–138.
- Nevo, A., Turner, J. L., and Williams, J. W. (2016). Usage-based pricing and demand for residential broadband. *Econometrica*, 84(2):411–443.
- O'Donoghue, T. and Rabin, M. (1999). Doing it now or later. *American Economic Review*, 89(1):103–124.
- Origin Energy (2020). Origin - good energy on the go (version 2.7.18)[mobile application software]. Retrieved from https://play.google.com/store/apps/details?id=com.originmobileapp&hl=en_AU&gl=US.
- Pope, D. G. and Schweitzer, M. E. (2011). Is tiger woods loss averse? persistent bias in the face of experience, competition, and high stakes. *American Economic Review*, 101(1):129–57.
- Rao, C. R., Rao, C. R., Statistiker, M., Rao, C. R., and Rao, C. R. (1973). *Linear statistical inference and its applications*, volume 2. Wiley New York.
- Rust, J. (1987). Optimal replacement of gmc bus engines: An empirical model of harold zurcher. *Econometrica: Journal of the Econometric Society*, pages 999–1033.
- Saez, E. (2010). Do taxpayers bunch at kink points? *American Economic Journal: Economic Policy*, 2(3):180–212.
- Sallee, J. M. (2014). Rational inattention and energy efficiency. *The Journal of Law and Economics*, 57(3):781–820.

- Sexton, S. (2015). Automatic bill payment and salience effects: Evidence from electricity consumption. *Review of Economics and Statistics*, 97(2):229–241.
- Simon, H. A. (1971). Designing organizations for an information rich world. In Greenberger, M., editor, *Computers, communications, and the public interest*, pages 37–72.
- Thakral, N. and Tô, L. T. (WP). Daily labor supply and adaptive reference points. Technical report, Working paper.
- Thaler, R. H. and Sunstein, C. R. (2009). *Nudge: Improving decisions about health, wealth, and happiness*. Penguin.
- The Australian Energy Market Commission (2017). *2017 Residential Electricity Price Trends*.
- Uetake, K. and Yang, N. (2018). Harnessing the small victories: Goal design strategies for a mobile calorie and weight loss tracking application. Available at SSRN 2928441.
- Werdiger, D. (2012). What is a bill shock and how can it be avoided? <http://www.billing.com.au/news/what-is-a-bill-shock-and-how-can-it-be-avoided/>.
- World Resource Institute (2018). Climate watch: Data for climate action. *Climate Watch*.
- Zizzo, D. J. (2010). Experimenter demand effects in economic experiments. *Experimental Economics*, 13(1):75–98.