

UNIVERSITY OF SOUTHAMPTON
FACULTY OF PHYSICAL AND APPLIED SCIENCES
Electronics and Computer Science

**People's Interaction with Future Autonomous Energy Systems in
Their Everyday Lives at Home**

by

Alper T. Alan

Thesis for the degree of Doctor of Philosophy

August 2016

UNIVERSITY OF SOUTHAMPTON

ABSTRACT

FACULTY OF PHYSICAL AND APPLIED SCIENCES
Electronics and Computer Science

Doctor of Philosophy

PEOPLE'S INTERACTION WITH FUTURE AUTONOMOUS ENERGY SYSTEMS
IN THEIR EVERYDAY LIVES AT HOME

by **Alper T. Alan**

Intelligent agents that sense and respond our continuing daily activities autonomously are becoming increasingly ubiquitous, and consequently transforming our lives. Domestic energy use accounts for a significant portion of national energy consumption in many countries, and is an important domain that intelligent agents may provide great benefit for us, by enabling a much more efficient energy utilisation. Whilst there are many algorithms developed for autonomous agents to assist people in managing their energy consumption at home, to date, there have been very few studies that examine human interaction with autonomy in the wild. Hence, there is a significant gap in our understanding of how people would react to and interact with autonomous agents in their everyday lives. This thesis aims to close this gap and help us to better understand how to design user interaction with future autonomous energy systems. To this end, we focus on people's perceptions of and interactions with two agent-based energy management systems that we designed and deployed based on envisioning future energy scenarios, and evaluated these systems through field studies. We represent implications for the design of future intelligent energy systems based on the results of our field studies.

The first system focuses on energy tariff switching. The decision of which energy tariff (i.e., energy pricing schema) to select is a challenging task for today's most households. Energy companies offer many different tariffs (e.g., standard, time of use and real-time tariffs) and it can be difficult to know which will be the most tailored to your consumption profile. Furthermore, the changes in the households' consumption and tariff rates increase the likelihood of ending up with a wrong tariff decision. To this end, we first focus on a future scenario where autonomous agents embedded in households have the ability to switch the energy providers daily, based on their offered rates and the households' consumption routines. To instantiate this envisioned scenario, we designed and developed two prototypes of a novel home energy management system called Tariff Agent, which monitors household energy consumption, as well as available energy tariffs, and therefore calculates the best tariff, and (optionally) automatically switches to it.

Both Tariff Agent prototypes offer flexible autonomy by which users can shift the system's level of autonomy in switching tariffs among three options: suggestion-only, semi-autonomous and fully autonomous, whenever they like. The first prototype was used by 10 UK households for 14 days. The findings from both quantitative and qualitative results of this first field study show that at least some people are ready to embrace software agents to manage their energy tariffs on their behalf as long as the agents reduce the hassle of tariff switching and maintain their budget. The results also indicate that although the users showed trust in Tariff Agent to control their tariff, they were still keen to monitor its performance. The second prototype was built based on the results of the first study and, differently from the first prototype, users are enabled to change the frequency of system reports that were previously sent once on each day of the study. To examine user interaction with the system for longer terms, the second study lasted 42 days and involved 12 UK households. The findings based on a thematic analysis show that flexible autonomy is a promising way to sustain users' engagement with smart systems, despite their occasional mistakes. The findings also suggest that users take responsibility of undesired outcomes of automated actions when delegation of autonomy can be adjusted flexibly.

The second system focuses on home heating. Home heating is a primary portion of energy expenses and therefore it is an important issue for residents. A number of smart thermostats have been introduced to customers to automate heating control on their behalf with the purpose of increasing the home's energy efficiency. However, none of these thermostats take into account energy prices that may vary based on residents' energy tariff. Hence, the second future energy scenario that we focus on envisions a smart thermostat that automates home heating control when energy price varies in real-time. To do so, we implemented three different smart thermostats that automate heating based on users' heating preferences and the real-time price variations. We evaluated our designs through a field study, where 30 UK households used our thermostats to heat their homes over a month. Our findings through thematic analysis show that the participants formed different understandings and expectations of our smart thermostat, and used its different features in various ways to effectively respond to real-time prices while maintaining their thermal comfort. Based on the findings, we present a number of design and research implications, specifically for designing future smart thermostats that will assist us in controlling home heating with real-time pricing, and for future intelligent autonomous energy systems.

Contents

Declaration of Authorship	xiii
Acknowledgements	xv
1 Introduction	1
1.1 Problem Statement	4
1.2 Requirements	5
1.2.1 Prototype Requirements	5
1.2.2 Evaluation Requirements	6
1.3 Research Objectives	7
1.4 Research Challenges	8
1.5 Research Contributions	10
1.6 Thesis Structure	12
2 Background	13
2.1 Human-Agent Interaction	13
2.2 Smart Grids and Real-Time Pricing	15
2.3 Autonomous Agents for Energy Management	16
2.4 Human-Computer Interaction for Energy Management	19
2.5 Models of Pro-environmental Behaviour	22
2.5.1 Rational Choice Models	22
2.5.2 Norm Activation Models	23
2.6 Motivation Methods for Pro-environmental Behaviour	23
2.6.1 Information	24
2.6.2 Goal-setting	24
2.6.3 Commitment	25
2.6.4 Feedback	25
2.6.5 Rewards and Penalties	26
2.7 Summary	26
3 Tariff Agent	29
3.1 Study Method	30
3.2 The First Study	30
3.2.1 Tariff Scenario and Tariff Specifications	31
3.2.2 Software Agent	32
3.2.3 Interaction Modalities of Tariff Agent	33
3.2.3.1 Home Page	33
3.2.3.2 Details Page	35

3.2.3.3	System Initiated Interactions via SMS	36
3.2.4	Findings	36
3.2.4.1	Engagement	37
3.2.4.2	Perception	39
3.2.4.3	Adoption	40
3.2.4.4	Accountability	42
3.2.5	Summary and Implications	43
3.3	The Second Study	44
3.3.1	Implemented Changes	45
3.3.2	Findings	46
3.3.2.1	Engagement	47
3.3.2.2	Perception	49
3.3.2.3	Adoption	51
3.3.2.4	Responsibility	53
3.4	Discussion	54
3.4.1	Interaction with Autonomy	54
3.4.2	Orientation to Smart Systems	56
3.4.3	Design Implications	57
3.5	Summary	58
4	Smart Thermo	61
4.1	The Study	61
4.2	Technology	63
4.3	Design Variations	64
4.3.1	Manual Thermostat	64
4.3.2	Direct Learning Thermostat	65
4.3.3	Indirect Learning Thermostat	67
4.4	Data Collection and Analysis	68
4.4.1	Interviews	69
4.4.2	Analysis	69
4.5	Findings	69
4.5.1	Overview of Quantitative Analysis	70
4.5.2	Orientation towards the Thermostat's Agency	70
4.5.3	Reactions to Different UI Features	71
4.5.4	Managing the Home Heating with Real-time Prices	71
4.5.4.1	Setpoint Preferences	71
4.5.4.2	Schedule	73
4.5.4.3	Boost	74
4.5.5	Mental Models of the Thermostat's Learning Feature	75
4.5.6	Balancing Cost and Thermal Comfort	76
4.5.7	Limitations of the Thermostat's Learning Model	77
4.6	Discussion	78
4.6.1	Designing a Thermostat for Real-Time Prices	78
4.6.2	Expectations from Smart Home Heating Systems	79
4.6.3	Studying Future Smart Energy Systems	80
4.7	Summary	81

5 Conclusions and Future Work	83
5.1 Conclusions	83
5.2 Future Work	85
A Real-Time Energy Prices	87
B Learning Heating Preferences	93
References	95

List of Figures

2.1	The Power-Aware Cord (left), The Eco-Eye (middle), The Wattson (right)	19
3.1	Home page.	34
3.2	Details page.	35
3.3	A text dialog under the semi-autonomous setting.	36
3.4	Overview of user activities. The rows listed under each user represent the total number of individual days that the user performed each activity. The last row shows the last day of user interaction.	37
3.5	Autonomy level changes over 14 days.	38
3.6	Page visits.	38
3.7	Acceptance of suggestions.	38
3.8	New home page.	45
3.9	New details page.	46
3.10	Overview of user activities. The rows listed under each user represent the total number of individual days that the user performed each activity. The last row shows the last day of user interaction.	47
3.11	Autonomy level changes over 42 days.	48
3.12	Page visits.	48
3.13	Acceptance of suggestions.	48
4.1	Overall system diagram.	63
4.2	Manual - home page.	64
4.3	Manual - schedule page.	65
4.4	Direct learning - home page.	66
4.5	Direct learning - settings page.	67
4.6	Indirect learning - home page.	68

List of Tables

3.1	Participants' profiles for the first study.	31
3.2	Tariffs in p/kWh.	32
3.3	Participants' profiles for the second study.	44
4.1	Participants' profiles.	62
4.2	Overall quantitative data analysis.	70

Declaration of Authorship

I, **Alper T. Alan**, declare that the thesis entitled *People's Interaction with Future Autonomous Energy Systems in Their Everyday Lives at Home* and the work presented in the thesis are both my own, and have been generated by me as the result of my own original research. I confirm that:

- this work was done wholly or mainly while in candidature for a research degree at this University;
- where any part of this thesis has previously been submitted for a degree or any other qualification at this University or any other institution, this has been clearly stated;
- where I have consulted the published work of others, this is always clearly attributed;
- where I have quoted from the work of others, the source is always given. With the exception of such quotations, this thesis is entirely my own work;
- I have acknowledged all main sources of help;
- where the thesis is based on work done by myself jointly with others, I have made clear exactly what was done by others and what I have contributed myself;
- none of this work has been published before submission

Signed:.....

Date:.....

Acknowledgements

First and foremost, I would like to thank my supervisors, Enrico Costanza and Sarvapali (Gopal) Ramchurn. They provided excellent advice and constant support throughout my PhD. Their commitment to science and endless passion to solve real-world problems have been by far the largest sources of motivation and inspiration for this thesis.

This work has also been possible through cooperation with excellent, friendly and helpful researchers. I am hugely appreciative to Prof. Nicholas R. Jennings, Prof. Sven Seuken, Prof. Tom Rodden, Joel Fischer and Mike Shann for their collaboration. I enjoyed and learnt a lot from working with them.

I would like to thank Enrico Gerding, my internal examiner, and Ya'akov (Kobi) Gal, my external examiner, for providing rigorous and excellent suggestions for improvement, and making the exam an enjoyable experience with their genuine friendliness.

The members of Agents, Interaction and Complexity (AIC) group have been great colleagues and friends. I will miss our house warmings, football and basketball matches, banquets, coffee breaks, and other socials. Thanks to everyone who made this lab a great environment to work in.

I gratefully acknowledge the funding sources, without which my PhD would not have been possible. The Ministry of National Education of Turkey provided my PhD scholarship. The studies presented in this thesis partially supported by the Engineering and Physical Sciences Research Council (EPSRC) ORCHID project (EP/I011587/1) and by the Swiss National Science Foundation (SNSF) project #147211.

Last, but not least, I am most grateful to my parents, my sister and Gizem for their unwavering blessings, love, affection, and trust in me. Their moral support and timely advice have been instrumental to my life and, in particular, the completion of this thesis.

Chapter 1

Introduction

The ways that people interact with computing systems are undergoing significant changes due to the ever accelerating progress of information technology ([Microsoft, 2008](#)). Our daily lives are increasingly becoming pervaded by interconnected ubiquitous devices such as laptops, tablets, mobile phones, wireless sensors and actuators. This creates a complex digital world around us that brims with vast amount of information ([Jennings et al., 2014](#)). However, people's capacity in monitoring, assessing and responding to all the information is limited. Thus, in order to complete tasks in such a complex digital environment, people have been utilising software agents, which are sophisticated computer programs that act for their users or other programs autonomously ([Wooldridge et al., 1995](#); [Jennings and Wooldridge, 1996](#)). For instance, today software agents are being used in various application domains including, but not limited to, personalised information management, electronic commerce, interface design, computer games, and the management of complex commercial and industrial processes ([Jennings et al., 2014](#)). In these domains, software agents not only provide users suggestions based on automated computations of information, but also learn users' preferences and complete difficult tasks autonomously.

On the one hand, autonomous operation may be desirable, or even essential, if we are willing to harness the opportunities offered by increasing amount of data. On the other hand, however, autonomous operation may not be the best choice due to noise and biases in real world data, the limited size of training datasets, and the discrepancies between computationally feasible models and complex real-life systems, which result in the operation of these “smart” autonomous systems being, at times, suboptimal or, in the worst case, detrimental. For example, recent work examining the real-world uptake of a smart thermostat highlighted how such errors are likely to cause users' frustration and may lead them to abandon the technology ([Yang and Newman, 2013](#)). It is therefore crucial that researchers and designers understand how to best design interfaces and interaction techniques that make the system status and operation clearly readable, and that allow its users to easily shift between autonomous and manual operation, a notion

known in the multi-agent systems community as “adjustable” autonomy (Scerri et al., 2003; Tambe et al., 2008).

To this end, the area of Human-Agent Interaction (HAI) has emerged to investigate the key design principles for establishing the interaction between humans and agents (Lewis, 1998). It focuses on the development of interaction design methods that determine how humans and agents interact with each other. To date, the majority of HAI research has explored human interaction with robots or virtually embodied agents (Holz et al., 2009). Furthermore, some of these studies were on specific domains that require considerable user training, such as aviation and military systems (Nourbakhsh et al., 2005; Hoc, 2000). However, there is a significant gap in our understanding of how we should design interactions with software agents, especially for the ones that might possibly intrude into our daily activities and have financial impacts on their non-expert users. We aim to explore and narrow this research gap in the domain of future energy infrastructures.

Future energy infrastructures, that are part of the vision called the smart grid, provide rich opportunities to explore interaction issues between humans and agents. The smart grid is defined by US Department of Energy (2008) as: *“A fully automated power delivery network that monitors and controls every customer and node, ensuring a two-way flow of electricity and information between the power plant and the appliance, and all points in between. Its distributed intelligence, coupled with broadband communications and automated control systems, enables real-time market transactions and seamless interfaces among people, buildings, industrial plants, generation facilities, and the electric network.”* This definition highlights that renewable energy resources, distributed intelligence, and autonomous energy systems will be major elements of our future energy grids. In this thesis, we focus on how autonomous energy systems may mediate user interaction with future energy grids. In particular, we are interested in the domestic setting, the potential of home energy technology to operate autonomously, and in people’s inclination, or resistance, to interact with such systems and to deliberately relinquish control to them.

The choice of home energy setting is driven by several factors. First, energy as an application is important in itself for its societal and economic implications (MacKay, 2009). Second, home energy systems provide an opportunity to study rich interactions with prototypes of future autonomous interactive systems “in the wild” since electricity is invisible (so much so that prior work aimed at materialising its representation, e.g., (Pierce and Paulos, 2010)), easy to measure with inexpensive and easy-to-install Internet-connected sensors, and its consumption is related to money in a way familiar to most users: through energy tariffs and bills. These characteristics make it possible to design and run field trials where financial experimental reward is linked to participants’ real electricity consumption and their use of prototypes, thus rendering the usefulness of the system more tangible to the participants. Indeed, exposing potential users to

functioning prototypes “in the wild” has long been recognised as critical for research of interactive systems (Rogers et al., 2007).

Now, the domestic energy domain has been an active area of research in Human-Computer Interaction (HCI) and Artificial Intelligence (AI) research disciplines. For instance, the prior work within the HCI community mostly focused on materialising the presentation of energy data (i.e., eco-feedback) to raise awareness and to promote desired behaviours, mostly for energy conservation (Pierce et al., 2008; Froehlich et al., 2010; Pierce and Paulos, 2012). In the same vein, the research area of AI developed various autonomous agents, scarce resource allocation mechanisms, intelligent coordination and control algorithms for optimising energy efficiency (Rogers et al., 2012). However, very few HCI and AI research have focused on the interaction issues between humans and agents, and the challenge of designing interactive autonomous systems. To our knowledge, there are yet no design guidelines derived from a real-world deployment for interactive autonomous energy systems that go beyond providing their users suggestions and automatically perform tasks on their behalf. For example, (Miller and Parasuraman, 2007) introduces 10 different levels of autonomy for an agent. However which of these levels would be the best for residential energy management, how and whether should the users share control with the agent, how and whether should the agent communicate its knowledge and planned actions, and how and whether should the agent request user input for clarification or notify the user of its actions are still unexplored questions.

Against this background, the main aim of this thesis is to address the challenge of designing future smart home energy applications and shed light on the interaction issues between humans and agents. In particular, our goal is to ascertain the principles for designing, implementing and evaluating interactive autonomous energy systems. In this thesis, we apply an existing HCI methodology to study human-agent interactions and to investigate the concept of autonomy within the domain of home energy management. In particular, we focus on the challenging problem of energy cost management given that energy prices dynamically change (e.g., real-time pricing) (we detail the problem in the next section), and how such a problem can be solved with the help of autonomous agents. To this end, based on envisioning future energy scenarios (Reeves, 2012), we designed and developed two agent-based energy systems: an autonomous tariff switching agent that helps its users in selecting energy tariffs at various levels of autonomy (see Chapter 3 for details) and a smart thermostat that automates home heating control to enable its users to cope with real-time prices (see Chapter 4 for details). We evaluated our systems with field experiments to explore user perception of and reaction to the agents in their everyday lives. Based on the quantitative and qualitative analysis of the field studies we provide novel design guidelines for developing future interactive autonomous energy systems.

The next section starts with the explanation of the tariff switching and the home heating problems, and later discusses why we need design guidelines for developing autonomous

energy systems that aims to assist people to cope with these two problems.

1.1 Problem Statement

The smart grid requires new energy and communication infrastructures that enable bidirectional flow of energy and information in energy distribution network. As a first step towards this development, smart meters have been installed in residential and commercial buildings in many countries ([The Department of Energy & Climate Change, 2012](#)). For example, the Smart Metering Initiative (SMI) driven forward by the UK Government establishes an obligation to install smart meters for all consumers by 2020 ([The Department of Energy & Climate Change, 2009](#)). Smart meters can measure electricity consumption over short intervals, typically every 30 minutes, and allow providers to offer time-based electricity pricing where the price for the electricity may change depending on the time of the day (e.g., time-of-use and dynamic pricing).

These time-based electricity tariffs have changed the consumers' role significantly and required them to be more active in the energy market. The consumers need to respond to variable prices in order to reduce their energy cost, either by reducing their energy use when the prices for the electricity are high, or by shifting their energy use to the cheaper periods. Although variable prices offer considerable savings, it is not easy enough for consumers to respond appropriately to time-based tariffs as mostly they are explicitly made confusing and complex ([Ramchurn et al., 2013](#)). In the long run, smart appliances are envisioned to respond autonomously to these price changes. However, throughout the transition period, consumers are likely to combat this complexity by themselves ([Rogers et al., 2012](#)).

To harness the benefit of variable energy prices, consumers need to continuously monitor their energy consumption and available energy prices. This appears to be a well-suited task for agents but not for humans when we consider that most people are not interested in dealing with energy. For example, a recent study has shown that people in the United States only spend two hours a year on average to search for better energy deals ([US Department of Energy, 2008](#)). Similarly, in United Kingdom, considering 26 million households, 7 out of 10 are on the wrong tariff, and between 40% and 60% of the households having the wrong tariff stick to the tariff they already have. The main reasons why people do not switch their tariffs are: the complexity of tariffs that are difficult to understand, the hassle of the tariff switching process where modification or cancellation fees may apply, and personal preferences, for example, one could prefer to stick to a tariff that is more expensive but more environmentally friendly.

To date, a number of comparison sites have been developed to help consumers find the best energy tariff. These services provide a list of comparison of various tariffs from different providers based on the user-provided estimate of how much energy will

be consumed monthly or annually. Although these services offer some suggestions and insights to consumers when choosing an energy tariff, their accuracy is highly depended on the consumption estimates provided by consumers. However, when we take into account varying consumption over different seasons and at different times during days, incorrect suggestions become inevitable. For example, predictions of annual energy bills for an average UK household can exceed £1,500 ([Frankcom, 2012](#)).

In this complex environment, interactive autonomous agents are required to support consumers so that they can actually benefit from varying energy prices by better understanding and controlling their energy use. For example, autonomous agents can be designed to serve households by learning their energy consumption profiles and switching their tariff to the cheapest option automatically. Another example could be a smart thermostat that can learn households temperature preferences over varying prices. Then it could automatically monitor and respond to the real-time prices on their behalf. By so doing, they could reduce the hassle of dealing with the complexity of dynamic energy prices whilst maintaining our budget. However, such automatic decisions might not be optimal all the time because of the uncertainties of users' energy consumption and the energy prices, and this may eventually lead users to abandon the use of agent technology. To render these autonomous systems practical, even under extreme uncertainties, there is a need for design guidelines derived from real-world deployments. We aim to introduce such design guidelines with this thesis. In the next section, we present the requirements for this research.

1.2 Requirements

This research does not focus on the problem of energy saving, and it goes beyond providing consumption feedback to explore human interaction with autonomous energy systems. To formulate novel design guidelines for autonomous energy systems for the domestic setting, a number of requirements need to be satisfied by this research. Firstly, we need fully functioning prototypes of an autonomous tariff switching agent and a smart thermostat that users will engage with. Secondly, we need to conduct field studies that will reveal how people interact with the agents in the real world. Therefore, we categorized the requirements of this research into two main parts: the prototype requirements and the evaluation requirements.

1.2.1 Prototype Requirements

The requirements need to be satisfied by our prototypes are as follows:

- **Configuration of Agency:** when and to what extent an autonomous agent-based system should perform actions on behalf of a human user should be configurable.

An autonomous energy management system should allow its users to individually balance user control and system autonomy due to the autonomy preferences that might vary for different users. For example, some users might favour ultimate system autonomy over full user control and authorize the agent to switch energy tariffs or heat their homes automatically based on varying energy prices, or prefer full user control.

- **Visibility of Agency:** in autonomous agent systems, agency should be visible in some degree to manage user expectations. However, the degree of visibility needs to be designed cautiously. An extreme visualisation of system agency (e.g., interface agents that act as direct intermediaries or guides) might evoke inflated user expectations and lead to user irritation. On the other hand, insufficient visualisation of agency might result in users being unable to observe system state and unaware of functioning agency. Therefore, the system agency should become visible in a lightweight way (i.e., not too intrusive) to inform the users of the changes of the system state, while keeping the complexity of the system infrastructure hidden from the users.
- **Control over Agency:** an autonomous agent-based system should be designed in a fashion that users can still have the control over the system's autonomous actions. To this end, the system's actions need to be coherent, intelligible and reliable. The consequences of the autonomous actions need to be clear to users and the users should be able to evaluate the risks associated with the autonomous actions. Furthermore, the users should be able to easily override the system's autonomous actions. In summary, the users should be able to understand what the agent is doing, why it is doing so, what it is going to do next and what is going to happen if the agent does it; and eventually the user should have the ultimate control over agency and should be able to direct the agent by overriding its autonomous actions.

1.2.2 Evaluation Requirements

The requirements need to be satisfied by the evaluation of an autonomous energy system through a field study are presented as follows:

- **Ecological validity:** it is important to conduct a field study to find out how people adopt and use a prototype in their everyday lives. However, for results of such a study to be meaningful it is important to provide a high level of realism, or in more formal terms a high degree of ecological validity. For example, in our case, to let participants experience the situation of an autonomous energy system affecting their budget, the system should use their real electricity consumption

data and should include monetary incentives based on their performance to mimic energy pricing.

- **Participant Diversity:** for the field study to be conducted, we require a number of participants who will use the system at their homes. However, the determination of the number of participants is challenging since the cost (time and money) of running such experiments grows with the number of users. The participants also need to have different demographics and cultural backgrounds so that the more representative findings can be obtained.
- **Quantitative and Qualitative Methods:** agent systems, so far, have been mostly evaluated through quantitative analysis. Although quantitative results are important means to support a hypotheses, it is hard to describe users' behaviour and activities only with quantitative data. For example, user interactions could be logged as quantitative data for seeing how many times a user visited a page but it is hard to reveal what the reasons were behind the visits of the page. Therefore, in addition to quantitative data, to reveal how users perceived our prototypes and reacted to them, we need qualitative accounts that can be collected through interviews.

Next, we present the research objectives of this thesis.

1.3 Research Objectives

Traditional approaches for defining the relationship between humans and agents (typically defined by the level of autonomy attributed to the agents) assume that humans are somehow more knowledgeable about the task and therefore more apt at defining when the agent should autonomously act and when it should request manual operation. This is typically referred to as adjustable autonomy, and it has been mostly used by the state-of-the-art technologies where humans rely on autonomous systems to complete complex tasks. However, in many cases, the points at which guidance needs to be requested may not be defined. Moreover, humans may instead be guided by agents to complement tasks undertaken by the agents. Thus, the locus of control may change between humans and agents at any point in time. Such interactions put humans and agents on the same level and are, hence, defined as realisations of flexible autonomy. Against this background, the aim of this research is to improve the state-of-the-art interaction methods for human-agent interaction in order to establish more sustainable relationship between human and agents in the specific domain of domestic energy consumption.

The ultimate objective of our research is as follows:

To develop design guidelines for autonomous energy systems that will be used by non-expert users in the context of home energy management. The guidelines should lead

to the development of agent-based systems in which autonomous operations are made configurable, visible, and controllable. The design of agents should provide lightweight interaction methods that are not too intrusive so as not to annoy the users in the long term, but that are sufficiently informative for building the users' trust in the agents.

1.4 Research Challenges

From the problem and requirements stated above there are research challenges that must be addressed. Here we discuss four key challenges facing the development of autonomous agents for real-world use.

- **Intelligibility of Agents:** The main driver behind the design of autonomous agents is the need to mitigate complexity of intricate tasks whilst ensuring that the automation is barely noticeable to its users (Klein et al., 2004). Meanwhile, it is important to design agents such that they present their status and intentions in an obvious way to users. In order for agents to be perceived as useful and reliable, their autonomous actions need to be adequately understandable and predictable. These two design principles contradict with each other, and this creates a design challenge: how to find the right balance between reducing the visibility of automation and increasing its intelligibility. For instance, Woods and Sarter (2000) have shown that high levels of automation in aircraft systems could lead to situations in which human pilots are not sufficiently aware of what the automation is currently doing, why it is doing it, and what it will do next. Norman (2002) describes mental and conceptual models in his book. Mental models are representations of a system that users create in their mind through interacting with the system (Weinschenk, 2011), and conceptual models are the actual models that designers utilise while developing a system. The system's design and interfaces can play a pivotal role in delivering the right conceptual model and therefore improving the intelligibility of agents. However, designing a legitimate conceptual model for autonomous systems to convey the right mental model to end-users is a significant challenge. In particular, excessive information presented in the interfaces, which communicate with the conceptual model, could overwhelm the users. While the presentation of insufficient information could complicate the intelligibility of agents' actions.
- **Agents' Interaction Frequency:** Proactiveness is the ability of agents to take initiative ahead of anticipated situations rather than reacting simply in response to their environment (Wooldridge et al., 1995), and this is a key element of agent autonomy (Nwana, 1996). A major challenge to design proactive agents is to generate accurate predictions for the relevant future, for example, future energy consumption or prices. However, the accuracy of predictions can be influenced by a number of different factors such as the prediction method used and the data

sensed. In particular, noisy data gathered from a sensing system may lead an agent to take faulty initiative. For instance, missing energy consumption data may cause an incorrect consumption prediction that is significantly lower than the actual amount. This erroneous prediction of consumption may consequently lead an agent to calculate an incorrect prediction of cost and therefore result in sub-optimal decisions. To increase the prediction accuracy and to reduce uncertainty in predictions, uncertainty reduction theory (Berger and Calabrese, 1975) could be utilized to develop interactive agents that humans can communicate with. For example, a consumer could inform an agent that her energy consumption for the next day's will be more than the predicted amount as the consumer has a better idea of what appliances (e.g., washing machine) will be used then. However, this communication might require considerable level of human interaction, especially for a consumer with a highly varying consumption values. Considering the fact that humans are not so willing to communicate with agents (Rodden et al., 2013), it is challenging to decide how often agents should request for human input. The agents should not become too intrusive with too many requests and they could still continue operating thoroughly.

- **Human Trust in Agents:** People are mostly reluctant to delegate their tasks to agent-based systems, although they might readily trust simple deterministic mechanisms that are transparent in their designs (Bradshaw et al., 2004). To help people to adopt autonomous energy management systems into their everyday lives, it is crucial to establish human trust in agents that hold the promise of improved quality of life. Human trust in agents is an essential factor to build long term engagement between humans and agents. People need to be convinced that agents are useful and dependable for critical tasks through showing that agents hold predictability and resemblance in their actions for certain circumstances. However, this challenge runs counter to the principle of making agents more adaptable since they may become less predictable due to the adaptation (Klein et al., 2004). Moreover, agents might cause harm more than good due to environmental uncertainties or design faults. For instance, an autonomous energy management system might lead a consumer to waste considerable amount of money because of faulty consumption and cost predictions that result from the uncertainties in demand and supply. Therefore, users must be able to evaluate the rewards and the potential risks of cooperating with such autonomous systems in order to trust them.
- **Flexible Autonomy:** Inadequate level of complexity and autonomy of agents can result in undesired consequences (Klein et al., 2004). To address this concern, a number of techniques have been developed to ensure that the autonomy of agents can be dynamically changed (Christoffersen and Woods, 2002; Myers and Morley, 2003). In those agent systems, there are mostly control policies defined to enable users to dynamically regulate the system's behaviour. Users are

able to define limits on autonomous behaviour according to their evaluation of an agent's proficiency. By so doing, undesirable outcomes can be prevented by managing the level of autonomy attributed to agents and making agents work with corrective measures (Klein et al., 2004). Furthermore, the capability to determine autonomy dynamically means that poorly performing agents can be immediately degraded to prevent undesired consequences (Bradshaw et al., 2004). However, the specification of the control policies can sometimes be cumbersome as the capabilities of agents and humans can vary with respect to a given context. The level of autonomy given to the agents can significantly influence the performance of tasks. However, selecting the appropriate degree of autonomy that determines what tasks will be delegated to agents on behalf of humans is quite challenging as low autonomy level could reduce the performance at dealing with any complex data while high autonomy level could limit human control and lead to undesired consequences. The challenge for this study is to decide which levels of autonomy to offer non-expert users in the context of home energy management.

The following section details our research contributions.

1.5 Research Contributions

In this thesis, we present the outputs of interdisciplinary research combining techniques and principles from HCI and AI, specifically by prototyping and deploying future autonomous smart grid applications and by evaluating them in the wild. Our key contributions can be summarised as follows:

- We present findings from two field studies where two different prototypes for automating energy tariff-switching were designed, implemented and evaluated in the wild with households having various lifestyles. Both prototypes offer flexible autonomy by which users can shift the system's level of autonomy among three options: suggestion-only, semi-autonomous and fully autonomous, whenever they like. We empirically demonstrate that flexible autonomy is a promising approach to sustain user interaction with smart energy systems. We present households' preferences over different levels of system autonomy for both short and longer terms. We show that users take responsibility for any undesired outcomes of automated actions when delegation of autonomy can be changed flexibly. We then provide novel design guidelines derived from real-world experiments for developing autonomous energy systems with flexible autonomy.
- We introduce the first smart thermostat study given real-time prices. The thermostat allows its users to control their home heating with real-time prices, and

automates the heating based on the temperature preferences of the users over different energy price levels. We evaluate our system with 30 UK households in a four-week in-situ study. We show that our participants formed different understandings and expectations of the system, and used it in various ways to effectively respond to real-time prices while maintaining their thermal comfort. Based on the findings, we provide design implications for developing a smart thermostat that autonomously responds to real-time prices on its users behalf.

The research presented in this thesis was also published in the following two full papers at international conferences:

Alan A, Costanza E, Fischer J, Ramchurn S, Rodden T, Jennings, N (2014). A field study of human-agent interaction for electricity tariff switching. In, Proceedings of the 13th International Conference on Autonomous Agents and Multi-Agent Systems, Paris, France, 965-972.

Alan A, Shann M, Costanza E, Ramchurn S, Seuken S (2016). It is too hot: an in-situ study of three designs for heating. In, Proceedings of the SIGCHI Conference on Human Factors in Computing Systems, San Jose, United States.

Moreover, the second study presented in Chapter 3 was presented at a workshop:

Alan A, Costanza E, Ramchurn S, Fischer J, Rodden T, Jennings, N (2015). Managing energy tariffs with agents: a field study of a future smart energy system at home. In, Adjunct Proceedings of the 2015 ACM International Joint Conference on Pervasive and Ubiquitous Computing and Proceedings of the 2015 ACM International Symposium on Wearable Computers, Osaka, Japan, 1551-1558.

Finally, a journal paper about the work presented in Chapter 3 is currently under review:

Alan A, Costanza E, Ramchurn S, Fischer J, Rodden T, Jennings, N (2015). Tariff Agent: interacting with a future smart energy system at home. In, ACM Transactions on Computer-Human Interaction (under review).

We next describe the structure of this thesis.

1.6 Thesis Structure

The remaining chapters of this thesis are structured as follows:

Chapter 2 provides the essential background information that is required to contextualise the work presented in this thesis. We start by outlining a brief review of human-agent interaction studies. Next, we discuss the smart grid vision and the current changes undergoing in global energy market to motivate studying human-agent interaction in energy domain. Then, we represent prior energy related literature from AI and HCI as the work we present in this thesis lies at the boundary of these two research areas. Finally, we review models and motivation techniques for proenvironmental behavior, which are necessary knowledge for designing energy management systems.

Chapter 3 introduces Tariff Agent, a novel agent-based tariff switching application. In this chapter, we first detail its major components: tariff specifications, software agent and interaction modalities. Next we present a field evaluation of the first Tariff Agent version and show how participants perceived and reacted to our system. Later we introduce the changes we implemented to Tariff Agent, based on the results of the first study. Then we present a field evaluation of the new version, and discuss the findings it revealed. Lastly, we discuss the results of the two field studies, and provide key design implications for developing future autonomous systems with flexible autonomy.

Chapter 4 presents Smart Thermo, a novel autonomous thermostat application. We start this chapter by providing the overall system description and the thermostat's design variations. Next, we present a field study where different households used each thermostat design to control their actual home heating over a month. We then state the findings of the study and discuss their implications for the design of future smart heating systems.

Chapter 5 gives a final summary of our research represented in this thesis. Firstly, we discuss the conclusions that can be drawn from each chapter. Later we outline future extensions of the work presented in this thesis, with specific attention paid to the potential for further user interaction design with autonomous agents.

Chapter 2

Background

Interdisciplinary knowledge and research is vital to facilitate a better understanding of how user interactions can be designed for smart technologies. The work we represent in this thesis lies at the boundary of AI and HCI. Hence, in this chapter, we present the background from both sub-disciplines of computing, which is required to understand the approach for developing interactive autonomous agents for home energy management. This chapter contains six main sections. In Section 2.1, we introduce an overview of human-agent interaction studies. Section 2.2 discusses the smart grid vision with the current implementations undergoing. Next, in Section 2.3, we represent the literature of agent technologies for energy management domain. In the same vein, we show human-computer interaction studies on energy management in Section 2.4. Finally, in Sections 2.5 and 2.6, we review some behaviour models and motivation techniques that we believe are crucial to understand before designing any interactive system.

2.1 Human-Agent Interaction

Maes et al. (1997) presented a debate on whether users should retain complete control of their interaction with interfaces or instead delegate some form of control to software agents that act on their behalf. In the debate, Ben Shneiderman emphasizes the importance of direct manipulation that is a human-interaction way which continuously presents the data of interest and offers reversible and incremental actions with rapid feedback. He expresses his concerns about the responsibility of the delegated actions. Therefore he advocates the idea that users need to comprehend the system which has to be predictable, and they need to be in charge of full control to take responsibility of their actions.

On the other hand, Pattie Maes (Maes et al., 1997) states that direct manipulation will eventually be unable to represent the computer environments that are becoming more

and more complex. She also highlights the issue that users are not highly computer-trained. Hence, she asserts that users have to delegate certain tasks, to some extent, to agents that can act on their behalf or at least provide suggestions based on their preferences. Moreover, she indicates that software agents are not alternative to direct manipulation but is a powerful complementary technology. She suggests that a well-designed autonomous system must still enable users to bypass the agents and perform actions manually when they want, and the system should improve users' trust in the agents by helping users to understand what the agent does.

Indeed, the potential of agents to mitigate the increasing complexity of computer environments has been recognized by HCI community, and increasingly applied to provide intelligent services([Shenghua, 2010](#)). For example, [Baker et al. \(2009\)](#) showed how agent-based systems can be applied in the domestic setting for improving energy efficiency. In the study consumers were enabled to interact with each home appliance (agents) to monitor and manage their energy consumption. Similarly, [Banerjee et al. \(2011\)](#) demonstrated how agents can help users to make better use of the renewable energy generated by providing suggestions. Understandably, by the development of such agent-based home energy systems, the agents will be more involved in consumers' everyday lives. Therefore, users will have more interactions with the agents at home ([Rodden et al., 2013](#)).

[Rodden et al. \(2013\)](#) investigated users' interactions with software agents embedded within future energy systems, and shows why it is crucial to examine human-agent interaction. In the study, animated sketches were used to convey the idea of software agents that are embedded into homes and have the functionality of electricity monitoring, switching energy provider and controlling appliances. The results of the study indicate that even though their participants showed a willingness to embrace the agents to assist in dealing with complex energy systems, there was a notable lack of trust in energy companies among the participants, which may eventually influence the consumer trust in software agents, as these agents might be provided by the companies.

So far, very few studies in the AI community have shown interest in the social aspects of the use of agents, and rather mainly concentrated on technical aspects such as algorithms, communication languages, ontologies, agent-oriented programming and verification of agent-based autonomous systems. However, understanding the interactions between users and agents is fundamental to the design of agent-based systems that people will adopt. An agent-based system with a poor design could lead to frustrations and inconveniences for the users that interact with the system. For instance, intrusive agent suggestions, agent failures or agents that do not fit human needs may reduce the usability of such agent-based systems and cause the avoidance of agent technology ([Shenghua, 2010](#)).

Approaches to human-agent interaction include different research topics (Bradshaw et al., 2011). The main research topics that have been formed by these approaches as some also stated in (Bradshaw et al., 2011) include interface agents and assistants (Clancey, 2004), adjustable autonomy (Scerri et al., 2003; Tambe et al., 2008), mixed-initiative systems (Allen et al., 1999; Horvitz, 1999; Zimmerman et al., 2009), human-agent teamwork (Sierhuis et al., 2003; Kamar, 2014), negotiation and repeated interactions (Gal et al., 2011; Peled et al., 2015; Rosenfeld et al., 2015; Azaria et al., 2015), and collaboration theory (Rich et al., 2001). Moreover, some recent studies (Traum and Rickel, 2002; Goodrich and Schultz, 2007; Murphy and Schreckenghost, 2013; Sauppé and Mutlu, 2015) focus on human interaction with robots or virtually embodied agents, and provide important insights about interacting with them. However, these studies neglect how software agents should be integrated in the infrastructure of smart applications and how they should communicate with their users.

In the next section, we present the vision of the smart grid and real-time energy pricing to better understand why the energy is a suitable domain for agent technology.

2.2 Smart Grids and Real-Time Pricing

The vision of smart grid is being adopted on a global scale, for instance in the UK where a roll-out of smart energy meters to all households is planned to be completed by 2020 (The Department of Energy & Climate Change, 2013). Energy distribution grids in which energy is simply delivered from suppliers to consumers are becoming more and more intelligent by sensing and responding to consumer demand. Meanwhile, energy generation sector is undergoing significant changes with penetration of renewables, including wind turbines and solar farms (Simm et al., 2015). However, energy production with the renewables fluctuates depending on the weather. This variability of energy generation puts more strain on the energy market to balance supply with demand, and increases the burden of dealing with peak demand that can lead to power outages. Peak demand further reduces the efficiency of current grids given the large capital investments needed to increase energy generation capacity to meet short peaks in demand (MacKay, 2009).

Moreover, the advent of electric vehicles (EVs), such as the Nissan Leaf and Chevy Volt, will change the energy requirements of transportation from fossil fuels to electricity (US Department of Energy, 2003). The batteries of these electric vehicles will introduce considerable additional load on the distribution grid, since an EV battery needs to charge quickly with a significant amount of energy to guarantee reasonable distances. In particular, an EV battery may be charged with 32 kWh of energy in few hours to ensure the range of around 100 miles, while a typical household may consume between 20 to 50 kWh of energy per day (Green et al., 2011). Thus, the requirement of the total

energy will be boosted by these vehicles. Furthermore, new peaks could be caused by these vehicles as their demand is likely to be concentrated on particular periods of the day. For instance, a local distribution network might become congested when supplying the aggregate demand when all the EVs in the local area are charged at the same time([Ramchurn et al., 2012](#)). Consequently, the burden of dealing with peak demand will be increased.

To overcome the issue of peak demand, utility companies and government agencies have been promoting demand response policies such as dynamic or real-time pricing where energy prices vary over short time intervals, typically hourly ([Barbose et al., 2004](#)). The aim of these pricing regimes is to incentivise consumers to reduce their consumption at peak periods where the prices are presumably higher ([US Department of Energy, 2006](#)). Real-time pricing has the most potential for consumers to benefit as consumers can obtain considerable monetary savings by shifting their loads to off-peak periods. However, it also provides the highest risk compared to less dynamic variations such as time-of-use pricing ([Faruqui and Palmer, 2011](#)). There has been early research and tested pilot programs have reported successes and opportunities for dynamic pricing ([Heberlein and Warriner, 1983](#)). A recent study with approximately 700 households in Chicago reports that consumers reduced their loads at peak periods due to price increases ([Allcott, 2011](#)). Research that focus more on how consumers sense real-time pricing in reality are scarce and appear to favour more simple pricing models ([Dütschke and Paetz, 2013](#)).

The dynamic or real-time pricing significantly change the role of consumers and offer lower energy bills. However, consumers are required to continuously monitor the changing prices to take the advantage of them. While this challenge might be a very daunting task for people, it would be a well-suited task for a smart system that can learn our preferences, monitor the prices, and respond to them autonomously. In the following section we discuss how autonomous agents can help people dealing with complex energy tariffs (e.g., real-time pricing).

2.3 Autonomous Agents for Energy Management

Autonomous agents are intelligent computing entities that can automatically operate on their users' behalf ([Wooldridge et al., 1995](#)). Although the definition of agents can vary, there are three commonly agreed fundamental attributes of an agent. These three main characteristics are as follows:

- Reactivity: ability to react to the changes of its environment.
- Pro-activeness: capability to make plan to achieve pre-designed goals.

- Social ability: agents could negotiate, cooperate or compete through communicating with each other or possibly with humans.

Agents can vary in terms of their intelligence level (Ferber, 1999). For instance, simple agents, also known as reactive agents, only perform basic reactions to stimuli that they monitor through some sensors. The use of these agents is especially prevalent when fast response time is vital (e.g., electric protection). However, these simple agents do not fit in complex tasks, which require more intellectual capabilities. More sophisticated agents such as intelligent agents or learning agents can achieve their goals through using their available resources and skills, or gain new knowledge about their environment through observations. In case of energy management, the most significant strength of agents is the ability of them to interact with each other. They can coordinate their activities and cooperate with each other so as to reach a common objective. By so doing, they could generate a form of distributed intelligence and act collectively (Wooldridge et al., 1995).

In recent years, there have been many research studies that use agent technology for solving energy related problems. In particular, agent-based technology has been used in demand side management in order to flatten the peak residential electricity demand that increases generation costs in terms of both monetary and environmental. To this end, there are many agent-based algorithms proposed for regulating home electricity demand (Bakker et al., 2010; Keshav and Rosenberg, 2010; Schülke et al., 2010; Bar-Noy et al., 2008). Although these studies used a similar system architecture including a communication network, a list of smart meters and programmable switches, their approaches and objectives were different, such as reducing total consumption and reducing costs based on dynamic pricing, or shifting loads or using batteries in order to match renewable generation. Similarly, SmartCap (Barker et al., 2012) focusses on residential loads and divides them into two main folds: interactive loads and background loads. Interactive loads generate the majority of loads in houses and they have little scheduling flexibility (lights, TVs, microwaves, etc.). On the other hand, a substantial part of home electricity consumption derives from background loads with some limited scheduling flexibility (air conditioners, heaters, freezer, refrigerators, dehumidifiers, etc.). Unlike the previous three works, SmartCap focuses on scheduling of background loads. Thus, it does not influence the comfort of occupants through inactivating or scheduling interactive loads. They achieve the decrease of average deviation from the mean power by over 20%, where the deviation is at least 400 watts.

Numerous other studies exploit agent technology to overcome various challenges in the energy domain (Alam et al., 2013; Vytelingum et al., 2011; Truong et al., 2013). Alam et al. (2013) take an agent-based approach to coordinate energy exchanges between households that are located off-grid and equipped with renewable energy generators and electric batteries in order to reduce the battery use and improve energy efficiency. To enable electricity consumers to obtain savings on their energy bills and maximise

social welfare, [Vytelingum et al. \(2011\)](#) introduce a novel agent-based micro-storage management method where individually-owned micro-storages are managed according to some learning strategies that adapt to market conditions. [Truong et al. \(2013\)](#) model the users' everyday routines and link the use of different appliances so as to predict future multiple appliance usage and, based on the predictions, to suggest the best time to run appliances when it is more beneficial in terms of saving money and reducing carbon emissions.

Furthermore it is possible to find numerous models and algorithms developed with different approaches for energy efficient heating. Some approaches focused on the models of the environment (e.g., the weather) to create an efficient heating schedule ([Yu et al., 2013](#); [Oldewurtel et al., 2010](#)). Other approaches used motion sensing to detect people's presence, and control the heating based on the occupancy models of buildings ([Scott et al., 2011](#); [Lu et al., 2010](#); [Panagopoulos et al., 2015](#)). [Shann and Seuken \(2013\)](#) presented a learning algorithm that elicits users' preferred temperatures for different energy prices and creates a comfort-cost trade-off model for each user. [Lam et al. \(2014\)](#) introduced a thermal comfort model that updates based on the user's comfort feedbacks. However, all these studies have simulation based results. Therefore, we cannot know whether they will actually work when they are deployed as a real world application.

Closer to one of our prototypes Tariff Agent, [Ramchurn et al. \(2013\)](#) introduce an advanced agent-based platform AgentSwitch that aims to solve the tariff selection problem encountered by residential electricity consumers. It included novel energy usage prediction and appliance disaggregation algorithms in order to be able to suggest that users shift their deferrable loads to off-peak times for increasing savings. Also, it provides a collective energy purchasing mechanism so that users can benefit from group discounts. The mechanism relies on the Shapley value, which ensures that the discounts are fairly shared by the users purchasing energy collectively. Furthermore, a novel provenance-tracking service is provided by the platform which aims to distribute accountability among individual system components and increase reliability of suggestions. AgentSwitch is a significant example of how novel agent technologies can be incorporated to solve energy-related real problems. In a usability lab study of AgentSwitch, [Fischer et al. \(2013\)](#) offer design recommendations for personalized energy-related recommender systems and underlines the potential for semi-autonomous systems with the challenge of balancing user control and autonomy flexibly.

Next, we explore HCI studies to understand how user interactions are designed so far to convey energy-related information, and to mediate human interaction with autonomous systems.



Figure 2.1: The Power-Aware Cord (left), The Eco-Eye (middle), The Wattson (right)

2.4 Human-Computer Interaction for Energy Management

Reduction in domestic energy demand at peak load periods would allow better integration of renewable energy sources and mitigate the growing generation capacity need caused by the peaks. A considerable increase in efficiency of electricity consumption and a significant decrease in the peak demand can be attained through informed consumers with energy management systems in which decisions are taken by agents or individual households. However, for such systems to succeed, it is crucial that actual power usage or energy related information is easily accessible and understandable for consumers.

The visualisations of the energy data play a significant role for the understandability of how the energy is being consumed, as mostly people are not familiar and comfortable with the energy domain. The main issues in transmission of this information are customers' understanding of energy concepts, convenient visualisations for conveying this information and appropriate concrete pattern for the visualisations (Monigatti et al., 2010).

To address these issues, HCI and Ubiquitous Computing researchers have been mainly focused on the development of energy feedback systems, also known as eco-feedback technologies, where energy consumption activities are perceived and related information is provided as a feedback (Pierce et al., 2008). Eco-feedback technologies aim to mitigate the adverse consequences caused by the lack of awareness and understanding of the people about how their each energy consumption activity can affect the environment by a set of creative and innovative visualisation techniques (Froehlich et al., 2010). In addition, these interactive instruments underpin and expand positive attitudes to sustainable activities. The figures below illustrate some of the creative visualisation techniques.

The Power-Aware Cord displayed on the left in Figure 2.1 is a novel power strip which visualises electricity consumption instead of hiding it. Electricity usage is visualised through bright pulses and changing density of light. Although the Power-Aware Cord is limited in terms of possible interpretations of its visualised data, it is effective in

improving awareness of consumers. The Eco-Eye shown in the middle of Figure 2.1 is one of the commercial devices providing real-time feedback on the total domestic energy consumption. It provides feedback through plain numerical data. The Wattson displayed on the right side in Figure 2.1 is a household energy monitor which provides ambient display in addition to direct real-time feedback. Live energy consumption is demonstrated numerically with ambient lighting. The colour of light states whether energy consumption is high or average. Previous studies indicate that approximately 10% energy saving is feasible through the use of real-time feedback systems such as the Eco-Eye or the Wattson (Dobson and Griffin, 1992; Mountain, 2006).

Against this background, Strengers (2011) asserts that current eco-feedback systems consider householders as resource managers that measure the cost of their consumption and make rational and efficient decisions. However, this assumption is missing other everyday life factors such as social and cultural situations. Therefore, those eco-feedback systems are likely to only attract the attention of environmentally motivated people and have volatile effectiveness. In the work, several alternative design directions are identified. For example, as householders often cannot understand resource management units (e.g., kilowatts and tons), these could be expressed by visual analogies such as buckets of water, which are more common in everyday life. In this vein, Costanza et al. (2012) proposes an interactive energy consumption visualisation tool that enables householders to play with the graphical representations of their historical energy consumption data with the aim of assisting users to better understand their energy consumption by relating their consumption data to concrete activities.

Darby (2006) reviews several previous studies about energy feedback, and highlights the effectiveness of energy feedback systems in improving energy conservation and efficiency. However, there are some considerations stated by (Pierce et al., 2008) on the evaluation of the effectiveness of such systems. They argue that the effectiveness of energy feedback is usually only assessed according to measured reduction in energy consumption. Therefore, they point out that there is not enough research focusing the effects of feedback on behaviours, understandings, adoption and social relations of householders. For example, which behaviours result in energy savings, how people sense the existence of an eco-feedback system (e.g., pleasing, enjoyable, helpful or useless) and how the system influences social relations within the household are the missing aspects of the existing feedback studies.

To address this gap, Pierce et al. (2008) introduce a mobile application that semi-automatically traces and reveals information about users' transportation behaviour. The application provides feedback by personal ambient displays based on their sensed or self-reported transportation behaviours. The aim of the study is to encourage people to use eco-friendly transportation such as walking, biking, taking public transportation or carpooling. The representation of eco-friendly transportation is used jointly with the representation of other goals (e.g., saving money, saving natural habitats, preserving

species and getting exercise) that may attract users' attention and eventually may lead them to adopt environmentally friendly transportation behaviours. According to the results of the study, it is hard to say that behaviour change occurred for all participants, however the participants engaged with the application and started showing new behaviours such as going to work by walking rather than driving a personal car. Although the study examined a semi-autonomous feedback system from a different perspective, it does not investigate the challenges of user interactions with semi-autonomous or fully autonomous agents.

In fact, to date, HCI research community has showed minimal interest in studying the interaction issues between humans and agents within domestic energy environment. Very few studies have explored how these new technologies (agents) could be harnessed to develop new ways for people to interact with energy and whether they would be adopted by real users and integrated into their everyday lives.

[Yang and Newman \(2013\)](#) have examined the real-world uptake of a smart thermostat with 23 participants. They highlighted how sub-optimal decisions taken by a smart thermostat are likely to cause users' frustration and may lead them to abandon the technology. Their follow-up study ([Yang et al., 2014](#)) has investigated users' long term interactions with the smart thermostat. Their findings suggest that users' interactions faded over time and resulted in unrealised energy saving opportunities. They also propose that an alternative design (i.e., a mixed-initiative system) might improve the sustainability of user engagement and the system's usefulness.

[Bourgeois et al. \(2014\)](#) deployed energy-aware washing machines that provide users suggestions on when to do their laundry based on the availability of green energy. They studied various intervention techniques with 18 households for 8 months, and showed that proactive suggestions sent by a software agent via text messages are more effective than the agent's email interventions. Similarly, [Costanza et al. \(2014\)](#) proposed Agent B, a software agent that also helps users book their washing machine in a scenario where electricity prices change in every 15 minutes. In a field experiment, 10 participants used Agent B for one month. The results suggest that Agent B helped users shift their laundry in response to real-time prices. The study also highlights the important challenge of how to determine the design features that achieve an acceptable balance between utility and convenience for users to adopt agent technology in their everyday lives seamlessly.

To design an energy-related autonomous system that will be adopted by users into their everyday routines, it is important to understand how people form energy-related behaviours and what factors motivate them to form these behaviours. Therefore, in the next sections we discuss the models of pro-environmental behaviour and motivation techniques for promoting such behaviours.

2.5 Models of Pro-environmental Behaviour

Pro-environmental behaviour corresponds to behaviour that does not damage the environment, or even benefits the environment. For instance, lowering thermostat settings, reducing car usage, waste recycling and composting, purchasing energy-saving light bulbs or preferring energy efficient appliances are various types of pro-environmental behaviour (Steg and Vlek, 2009). Understanding the underlying motivations of people who exhibit environmentally friendly behaviour has been studied by different disciplines including Economics, Sociology and Psychology. There are numerous theoretical models developed in order to provide insights about people's pro-environmental behaviours. These models clearly influence the design of interactive systems since designers take some model of human behaviour as basis when attempting to solve a problem (Froehlich et al., 2010). The most commonly used models of pro-environmental behaviour in the literature can be categorized into two basic types: rational choice models and norm-activation models. This section presents these two main models based on the literature review presented in (Froehlich et al., 2010).

2.5.1 Rational Choice Models

A rational choice model assumes that people are self-interested and environmental behaviour is triggered by the evaluation of expected utility. Attitude models are the earliest and simplest rational choice models. Attitude models are based on three inter-related components: affective component, behaviour component and cognitive component. Briefly, attitude models assume that knowledge leads to feelings and emotions, which eventually result in pro-environmental behaviour. However, these models do not consider the other possible factors that might affect the behaviour. Hence, attitudes do not always match with actual behaviour. For instance, an early study showed that people who consider energy conservation as the most significant way to deal with energy crisis are not different from others in energy-conservation behaviours (Costanzo et al., 1986).

Another rational choice model is the model of responsible environmental behaviour which considers additional factors. Hines et al. (1987), who proposed this model, state that intention to act and situational conditions such as economic constraints and social pressures play a significant role in the determination of pro-environmental behaviour. Lastly, the rational-economic model relies on the idea that people will demonstrate environmentally responsible behaviours if these behaviours bring economic benefits. However, it is not always the case that people match behaviours with costs. The rational-economic model also disregards the impact of non-economic determinants such as comfort, habit and social norms (Yates and Aronson, 1983). Furthermore, slight price manipulations might not carry great importance for people and might not lead to significant behaviour

change, while dramatically changing the cost of resources (e.g. water, gas and electricity) could easily become an ethical issue (Steg and Vlek, 2009).

2.5.2 Norm Activation Models

Norm activation theory indicates that personal norms are the only determinants of prosocial behaviour (Schwartz, 1977). Personal norms are assumed to be strong feelings of moral obligations that direct people to engage in pro-social behaviour. According to Schwartz (1977), the personal norms emerge from two premises: the recognition of the consequences of one's behaviours and the acceptance of personal responsibility for those behaviours. For instance, in norm activation model, if someone is aware of the effects of her fuel consumption on the climate change problem and accept the responsibility of fuel consumption behaviour, then the person is likely to develop personal norm to reduce her fuel consumption. Norm activation models differ from rational choice models as they suggest that behaviour might be stimulated by altruistic norms and moral obligations might outshine the subjective perceptions of utility (Staats et al., 2004).

The value-belief-norm theory introduced by Stern (2000) extends the norm activation theory through establishing a more sophisticated relationship among values, beliefs, attitudes and norms. In this theory, pro-social attitudes and personal norms are suggested as significant determinants of pro-environmental behaviour. The theory postulates that pro-environmental behaviour is primarily gained by the awareness of consequences in the norm activation model. However, the degree of acceptance of pro-environmental behaviour is correlated with personal values. As an illustration, if a person possesses strong altruistic and biospheric values, the person is more likely to adopt pro-environmental behaviour. This acceptance is less likely when the person holds egoistic values (De Groot and Steg, 2007). Thus, environmentally desirable behaviours are activated not just by the attention paid to other people who might suffer from environmental damage but also according to the self and non-human entities (Steg and Vlek, 2009). In the next section, we discuss some of the techniques that have been used to motivate pro-environmental behaviour.

2.6 Motivation Methods for Pro-environmental Behaviour

Designs based on the models of pro-environmental behaviour on their own, are not sufficient for encouraging people to change their behaviour (Froehlich et al., 2010). Thus, there has been research on particular interventions which might motivate people to adopt environmentally responsible behaviour. In what follows, we individually detail some of the most widespread motivation techniques including information, goal-setting, comparison, commitment, feedback, rewards and penalties.

2.6.1 Information

Information is the instrument which has been mostly used to promote pro-environmental behaviour through media campaigns, leaflets or websites. The underlying idea is that better informed people will behave more environmentally friendly. However, representation of the information about the benefits of pro-environmental behaviours typically does not lead to permanent effect. In order to improve the effect of information, it must be easy to grasp, reliable, easy to remember and delivered in an attractive way in the right place at the right time (Froehlich et al., 2010).

Geller (1981) evaluated the impact of a workshop in which energy-saving information was provided. The workshop increased the concerns about energy crises, improved the knowledge about energy conservation and strengthen the intentions of people to adopt energy-saving behaviours. However, the evaluation showed that although information influenced underlying determinants of energy usage, it did not lead to behaviour changes. In the same vein, studies by Luyben (1982), Hutton and McNeill (1981), and Staats et al. (1996) evaluated mass media campaigns and indicated the similar result: an increase in knowledge but not a necessary increase in willingness to behave pro-environmentally.

2.6.2 Goal-setting

Goal-setting is a powerful way of motivation that works through a comparison of the present and an ideal future situation (Van Houwelingen and Van Raaij, 1989). Goals can be set by authorities or by households. Other motivation techniques such as feedback and commitment are often used in joint with goal-setting so as to indicate how households are performing relative to the goal, or in order to increase compliance to the goal (Becker, 1978). Locke and Latham (2002) indicates that goals influence behaviour through directing attention and effort toward goal-related activities and enhancing persistence. Additionally, Craig and McCann (1978) states that the established goals should be feasible.

Becker (1978) argues that a relatively difficult goal is more effective than a relatively easy goal. The study states that households received a relatively difficult goal (save 20% energy) and feedback about their performance conserved the most (15.1%), while the households who had been given a relatively easy goal (save 2% energy) showed that an easy goal is not effective at all. Similarly, Van Houwelingen and Van Raaij (1989) concluded that goal-setting supported by daily feedback on consumption reduced natural gas use by 12.3%. Furthermore, McCalley and Midden (2002) found that participants who received a goal and feedback saved more energy per washing trial than participants who received only feedback. No considerable difference had been found between participants who had set a goal themselves and those with an assigned goal.

2.6.3 Commitment

A commitment refers to an oral or written pledge or promise to behave in a certain way or achieve a specific goal (Abrahamse et al., 2005). A commitment, which plays an important role as a determinant of environmentally friendly behaviour, can be made to oneself or to others, in either case it might activate personal or social norms. The impact of commitment varies with the type of commitment that is made, the person or people to whom the commitment is conveyed and whether the commitment is made privately or publicly. Pallak and Cummings (1976) found that households who committed to declare their results publicly consumed 15% less natural gas and 20% less electricity than other groups. Similarly, Wang and Katzev (1990) stated that a signed pledge to recycle raised 40% of recycling compared to baseline data.

2.6.4 Feedback

Feedback is another general method that is applied to promote pro-environmental behaviour. Feedback consists of providing households information about their performance towards a specific goal (i.e., energy reduction). Becker (1978) showed the positive effect of feedback on performance in connection with residential energy conservation. A majority of studies related to the effect of feedback on pro-environmental behaviours has focused on domestic resource consumption including water, electricity and natural gas (Froehlich et al., 2010). Feedback can vary in terms of the frequency of feedback, the level of feedback and the content of feedback. The most common feedback frequencies are continuous feedback, daily feedback, weekly and monthly feedback. Whilst low-level feedback aims to change or improve a certain behaviour through providing detailed information about the behaviour, high-level feedback provides summation information to improve performance towards a more general goal. Besides including individual performance information, feedback can contain performance information of others which might evoke a social norm (Abrahamse et al., 2005).

Cook (1979) showed that households who received continuous feedback consumed 12% less electricity than a control group. Similarly, Bittle et al. (1979) indicated that households that received daily feedback saved an average of 4% on their electricity usage compared to baseline consumption. A more recent study by Völlink and Meertens (1999), used a combination of motivation techniques including weekly feedback, goal-setting and information (energy saving tips), and showed that households with the combination of interventions saved more energy than the control group. Siero et al. (1996) conducted a study of energy consumption at two units in a company and found that the unit who received comparative feedback saved more energy than the unit who was subject to individual feedback. However, Haakana et al. (1997) and Egan (1999) argue that, although comparisons appeal people's interest, they might not always result in behaviour change.

2.6.5 Rewards and Penalties

Rewards and penalties are both consequence interventions which are based on the idea that the application of positive or negative consequences will have impact on behaviour. In other words, pro-environmental behaviours will be more preferable when rewarding consequences offered with them, and non-environmental behaviours will become more avoidable when negative consequences are attached to them (Abrahamse et al., 2005). Rewards and penalties are usually monetary. For instance, encouraging households to invest in home insulation in order to save money on their heating bills, or promoting the use of energy-efficient home appliances so as to reduce overall energy consumption and cost. However, the introduced consequences need not always be monetary, they could be related to status or convenience. For example, private reserved parking spots for ride-sharing have been demonstrated to increase carpooling, and residential curbside pickup services has significantly increased recycling efforts through making it easy to discard wastes (Stern, 1999).

Slavin et al. (1981) used monetary rewards with feedback in their two studies, which were conducted to explore the influences of a group contingency for conservation on electricity use. In both studies, average electricity savings of 6-7% relative to baseline was achieved. It was shown that people could even respond to symbolic reward such as an acknowledgement of pro-environmental behaviour. Consolvo et al. (2008) found that even displaying an asterisk after the completion of a behaviour can lead a positive response. There are eco-feedback designs that utilise game-like rewards (e.g., scores, levels, etc.) to promote positive behaviours (Bang et al., 2007; Consolvo et al., 2008).

2.7 Summary

This chapter first introduced a brief literature of Human-Agent Interaction studies. Secondly, the smart grid vision and undergoing developments in the energy domain was discussed. Then, a detailed overview of the key AI and HCI research related to this thesis were presented to provide the necessary background. Later, this chapter highlighted two main models of pro-environmental behaviour: the rational choice model which states that pro-environmental behaviour is initiated by evaluation of expected benefits, and the norm-activation model that considers personal norms as fundamental determinants of pro-social behaviour. Although these two models are not complete guides to explore all human behaviour, they can be still utilized by interactive designs in order to uncover behaviour factors. For instance, a design based on the rational choice model might highlight the monetary benefit of an environmental behaviour. Whilst a design based on norm-activation model could stress the foreseen consequences of an pro-environmental behaviour in wildlife to promote an altruistic attitude. Both in Tariff Agent and Smart

Thermo, we exploited the rational choice model with the assumption that users would interact with the system to improve their utility (i.e., monetary saving).

Lastly, this chapter presented the key motivation methods used in literature to influence households' energy consumption behaviour. Namely, these predominant motivation techniques are information, goal setting, commitment, feedback and rewards and penalties. As it is stated in the discussion above, the effectiveness of these techniques varies. Also, most of the techniques were used in combinations. Therefore, it is difficult to indicate the exact influence of an individual technique. For instance, most often goal setting or commitment is used with feedback so that people who made a commitment or set a goal can monitor how they are performing and how they are close to their pledge or goal. In case any motivation methods will be used, it is essential to initially determine what behaviours a design is aiming to motivate, as each behaviour includes different complexities and nuances. Hence, user interaction designers need to consider the behaviours that they are planning to influence before creating any system. Both in Tariff Agent and Smart Thermo, we aim to encourage people to be more active in understanding their consumption and reacting to changing energy prices. To do so, we exploit the combination of motivation methods: information, feedback, rewards and penalties. Chapters [3](#) and [4](#) demonstrate the use of these motivation methods in detail.

Chapter 3

Tariff Agent

In this chapter, we aim to study how people interact with an autonomous system that can proactively make decisions, which may have financial consequences and impact on the daily routines of its owner. In order to obtain meaningful results from such a study, it is crucial to offer a high degree of ecological validity. Therefore we designed and carried out two field trials by deploying different versions of Tariff Agent to observe how people make use of the system as a part of their everyday life, for both short and longer terms.¹ The system development is inspired by the future scenario depicted in a previous work (Rodden et al., 2013), where autonomous agents are embedded in households have the ability to switch the energy providers based on their offered rates and the households' consumption routines in order to ensure that the households are on the best tariff. The concept of switching energy suppliers or tariffs is familiar to households in the UK, especially as a number of government media campaigns in recent years encouraged consumers to change often in order to save on energy bills². Domestic energy consumers in the UK can switch their tariff through websites or call centres, as often as they wish³, but the switching process may take time, on average around 17 days⁴. We condensed this actual scenario in order to be able to run field studies of limited duration. Therefore, in our field studies, we assume that every household can switch its electricity tariff every day, one day in advance, through our agent-based application.

¹Ethics approval references for the two field studies are ERGO-10369 and ERGO-11382.

²<https://www.gov.uk/government/publications/household-energy-savings-through-switching-supporting-evidence>

³Some contracts might be binding for a period of time (e.g., 12 or 24 months).

⁴<http://www.moneysupermarket.com/gas-and-electricity/switching-suppliers/>

3.1 Study Method

To recruit participants we utilised personal contacts and snowball sampling, with the constraint that all participants would be blind to the purpose of our research. The recruitment criteria for participants were to have a broadband Internet connection, basic knowledge of Internet use, and to live in flats or houses where off-the-shelf energy monitoring kits could be easily and safely installed without the intervention of a professional electrician. Participants were asked to install the kit on their own and make sure it runs throughout the study, but support was provided where necessary. After the installation of the monitoring kits, our agent-based web application was introduced to participants. In particular we emphasised that the system does not affect their actual energy tariffs and bills, yet their daily energy cost and therefore the monetary reward that they will receive at the end is calculated based on their actual energy consumption and the tariff they select on the system. We detail the use of monetary reward below.

To motivate participants to engage with and experience the use of an autonomous system that may affect them financially, we provided monetary incentives based on their study performance. This idea of using monetary incentives to mimic energy pricing was inspired from earlier studies (Slavin et al., 1981; Costanza et al., 2014), where participants were rewarded according to their study performance. At the beginning of our studies, a certain level of budget was allocated to each participant, and their daily consumption cost was deducted from this budget over the period of the trial. At the end of the studies participants received payments (in the form of a shopping voucher) according to the amount left on their budget. By so doing, we aimed to make their savings have a real and tangible impact, and therefore encourage participants to engage with the system.

To reveal users' orientation towards an autonomous energy system we performed both quantitative and qualitative analysis. Quantitative data was collected through automatically recorded system logs, documenting the interaction of users with the system (e.g., how many times participants visited a specific page or changed the system's autonomy level). However, such data alone does not provide enough information to understand users' behaviour, for example why they would opt for a certain autonomy level or visit a certain page more frequently. Therefore, we complemented it with semi-structured exit interviews. These interviews focused on participants' use, adoption and understanding of the system, and lasted between 20 and 30 minutes. In the interviews, we asked open questions mostly related to their actual experience with the deployed system.

3.2 The First Study

We recruited a total of 10 participants (5 female) for the first field study. The first field study was conducted for 12 days, where participants engaged with the initial version

Table 3.1: Participants' profiles for the first study.

PARTICIPANT	GENDER	AGE	OCCUPATION	OTHERS
Alisa	Female	20s	Lawyer	4 Adults
Claudia	Female	20s	Chemistry UG Stud.	3 Adults
Ender	Male	30s	Law PhD Stud.	1 Adult
Greta	Female	20s	HR Manager	3 Adults
Ivan	Male	20s	Comp. Sci. PhD Stud.	2 Adults
Louisa	Female	30s	Manager	1 Adult, 1 Child
Maria	Female	30s	Comp. Sci. PostDoc	1 Adult
Mehmet	Male	20s	PhD in English	None
Omar	Male	30s	Comp. Sci. PhD Stud.	1 Adult, 2 Children
Sinha	Male	30s	Manager	1 adult

of Tariff Agent. At the beginning of the study all participants were allocated a budget of £30 for spending during the study. Half of the participants were local professionals working in different sectors, while the other participants were members of the university (see Table 3.1 for detailed profiles).

In what follows, we first explain our tariff scenario and tariff specifications. In Section 3.2.2 we explain how the autonomous system works. Then we show interaction modalities of the system in Section 3.2.3. In Section 3.2.4, we present the analysis of quantitative and qualitative data.

3.2.1 Tariff Scenario and Tariff Specifications

We consider a daily electricity tariff switching problem so as to be able to create a field study of limited duration (as depicted in the previous section). We assume that every household can switch tariff every day one day in advance, and can select its tariff only from two types of suppliers. The first type purchases their energy from traditional fossil fuel power stations, therefore they have dependable production for fluctuating demand, and sell electricity at a constant rate (r_{std}), in our scenario 15 p/kWh (pence per kilowatt-hour). The second type of supplier sells electricity at a variable rate that may change every day. This type of supplier buys its day-ahead energy from a wind generator that has different production every day depending on the weather. It needs to buy extra energy from the real-time market to meet possible shortfalls in supply. Thus, it passes such costs to its consumers through offering two rates: a low rate r_{wind} is applied to the consumption covered by the wind generation, while a high rate r_{nowind} is applied to the rest of the consumption.

To make the tariff switching problem more challenging for the users and therefore encourage them to consider delegating the process to an agent, we introduced three different suppliers from the second type, named as A, B and C (see Table 3.2). Each of these

Table 3.2: Tariffs in p/kWh.

Tariff	r_{wind}	r_{nowind}	Risk Level
Variable-A	3 p/kWh	23 p/kWh	High
Variable-B	8 p/kWh	18 p/kWh	Medium
Variable-C	10 p/kWh	16 p/kWh	Low

three suppliers offers a tariff with a different combination of low rate and high rate, corresponding to a different level of risk for the user. Tariff Variable-A is considered high risk because the gap between the low rate and the high rate is highest (r_{wind} : 3 p/kWh, r_{nowind} : 23 p/kWh), Variable-B is medium risk (r_{wind} : 8 p/kWh, r_{nowind} : 18 p/kWh), and Variable-C is lower risk (r_{wind} : 10 p/kWh, r_{nowind} : 16 p/kWh). The standard rate lies between the r_{wind} and r_{nowind} rates offered by these three tariffs. Therefore, it is not easy to decide which tariff is the cheapest.

In particular, there are two types of uncertainty affecting the tariff selection decision. The first one is personal uncertainty in predicting the user's own consumption for the next day, for example consumption might change with unpredictable visits by friends for a dinner or running out of clean clothes and therefore needing to do a wash. The second uncertainty is environmental since the availability of wind energy is weather dependent and therefore, while hourly predictions may be reasonably good, day-ahead predictions are likely to be inaccurate (Truong et al., 2013).

3.2.2 Software Agent

In our tariff switching scenario, planning which tariff to select and when to change it is a well-suited task for a software agent since it is necessary to continuously monitor the changing consumption and wind energy to predict the best tariff.

Participants' energy consumption is monitored through off-the-shelf home energy monitoring devices (AlertMe⁵). These devices measure the total consumption of the household through a current clamp, and make the data available through an HTTP API. Predicting day-ahead usage accurately has been shown to be a challenging problem, and even with sophisticated machine learning techniques produce unsatisfying results (Truong et al., 2013). Therefore, we implemented a simple prediction method, which uses the previous day's consumption as a prediction for the next day's consumption. As detailed in Section 3.2.3.1, users can manually provide the agent with a more accurate prediction of their consumption for the following day through the web interface.

To simulate the availability of wind generation, we collected wind data from a weather forecast service for 28 days from regions where wind turbines are located in the UK. We collected the actual and predicted wind energy values, which are then calibrated

⁵www.alertme.com

for individual users, based on their average daily consumption. This calibration is important because each individual user of Tariff Agent may have different daily electricity consumption levels, so we need to calibrate the wind energy values to ensure users' consumption may, at times, be higher than the available wind energy, and, at other times, be within the available wind energy. Our aim is to introduce uncertainty in the system, so that the agent would at times provide correct suggestions, but at times incorrect ones because of incorrect wind predictions and mismatch between available wind energy and users' consumption. Essentially we shift the wind probability distribution so its average and range roughly match the probability distribution of the energy consumption of each participant.

Even though realised and predicted wind energy values are stored on the same platform, the realised values are never passed on to the tariff switching algorithm used by Tariff Agent. To select or suggest an energy tariff, the agent first computes the predicted day-ahead costs for all tariffs by using the predicted user consumption and predicted wind energy for the next day, and then chooses the cheapest tariff. The realised wind energy values are then only used for calculating the daily cost at the end of each day, based on the user's selected tariff and actual consumption.

3.2.3 Interaction Modalities of Tariff Agent

Interactions between users and our system occur through two mediums: they can interact with the system either through SMS or through a web site. The site includes two pages: home and details, which are described in what follows.

3.2.3.1 Home Page

The home page, illustrated in Figure 3.1, comprises three components: tariff, setting and budget. Through the **Tariff** component users can see the current tariff and manually select the tariff for the next day. The current tariff is displayed on the top, highlighted in green if it is the same as the one the system suggested, otherwise in orange, to emphasise the difference.

In the middle of the component, the predicted values for the user's consumption and wind generation on the next day are shown. Below the predictions, the four tariffs are listed, from the one predicted to be the cheapest to the most expensive. The suggested tariff (the cheapest) is marked as such through a text label. Users can select a tariff through a button and they can bring up a detailed description of each tariff, including the rates, by clicking on the 'information icon' next to it.

The predicted amount of energy consumption can be modified through radio buttons. Changes in the consumption prediction are immediately reflected in the estimated costs

Tariff

Your current tariff is **Fixed Tariff**. Tomorrow's tariff is **Fixed Tariff**. You can change it before 9pm today.

Predictions for Tomorrow:

- 23.4 kWh - A lot more than yesterday
- 18.7 kWh - More than yesterday
- 15.6 kWh - **Same as yesterday**
- 10.9 kWh - Less than yesterday
- 7.8 kWh - A lot less than yesterday

Estimated Consumption:

Estimated Wind Energy Generation: 6.7 kWh

Tariff	Cost	Best	Worst			
Fixed	£ 2.03	£ 2.03	£ 2.03	Selected	Suggested	i
Variable-C	£ 2.09	£ 1.56	£ 2.50	Select	i	
Variable-B	£ 2.14	£ 1.25	£ 2.81	Select	i	
Variable-A	£ 2.25	£ 0.47	£ 3.59	Select	i	

[Save into agent settings](#)

Setting

- Send me an SMS when tariff change is suggested.
- Automatically select best tariff and send confirmation.
- Automatically select best tariff without confirmation.

Budget

Available: £19.40 **Spent:** £10.60

[Account Book Details](#)

Figure 3.1: Home page.

for each tariff. The manually selected prediction can be confirmed, or “saved”, to the system by clicking a button. By so doing, users can understand how the system uses their predicted consumption to make a better choice on their behalf and therefore inspire confidence in the system.

Setting is the second component of the home page, and it allows users to select one of the following three autonomy levels:

- Suggestion-only: If the system detects that the current tariff is different from the one predicted to be the best for the next day, it sends an SMS suggesting a tariff change. Users can accept the suggestion by replying “Yes” via SMS.
- Semi-autonomous: The system automatically switches to the predicted best tariff and informs users of the change via SMS. If the users are not happy with the change they can go to the website and manually change the tariff there. This level is semi-autonomous in that it automatically switches tariffs, but it allows users to easily regain control.
- Fully autonomous: The system automatically switches to the predicted best tariff but does not inform the user of the change. This level is fully autonomous in that it completely offloads users of the burden of tariff switching.

The users who have selected either the semi-autonomous or fully autonomous option could access the website and modify the predicted consumption, which will automatically

TariffAgent Home									
Account Book Details									
Date	Predicted Cons. (kWh)	Actual Cons. (kWh)	Predicted Wind Energy (kWh)	Actual Wind Energy (kWh)	Agent Suggestion	Selected Tariff	Cost (£)	Saved/Lost (£)	
19-Sep	5.4	5.6	2.3	2.3	Fixed	Fixed	0.73	0.10	
18-Sep	6.1	5.8	2.6	2.3	Fixed	Fixed	0.75	0.12	
17-Sep	4.3	5.4	2.4	2.1	Variable-A	Variable-A	0.82	-0.12	
16-Sep	5.7	6.1	2.2	2.4	Fixed	Fixed	0.79	0.13	

Figure 3.2: Details page.

affect the system's selection of tariff. To select a tariff manually, the users need to change their setting to suggestion-only autonomy level.

The last component on the home page is the **Budget**, which displays how much was spent and how much remains of the budget allocated at the beginning of the study. It also provides a link to the other web page of the system, which is described in the next section.

3.2.3.2 Details Page

The details page, shown in Figure 3.2, provides historical information about the operation of Tariff Agent, with the aim of allowing users to evaluate their performance, and make the system accountable. In particular, for each past day the predicted and actual values for energy consumption and wind generation are shown, together with the suggested and actual tariff selection, the cost and the saving or loss incurred. The saving incurs when the cost of the selected tariff is the cheapest compared to other tariffs' cost, and the value of the saving is determined by comparing the selected tariff's cost with the most expensive tariff. On the other hand, the loss incurs when there was a better tariff than the selected one, which would cost less money for consuming the same amount of energy. The value of the loss is calculated by comparing the cost of the selected tariff and the cheapest tariff.

To facilitate the understanding of the information displayed, values in the table are colour-coded. Tariffs are displayed in green or red depending on whether the selection was optimal or sub-optimal. Consumption and wind generation predictions are shown in green when they turned out to be accurate (within 10%) and the resulting tariff suggestion is optimal. They are shown in red when they are inaccurate (outside 10%) compared to the realised values and the resulting tariff suggestion is suboptimal. They are shown in orange if the predictions are off (outside 10%), but the resulting tariff suggestion was optimal (for example in the case that one error compensated for the other).

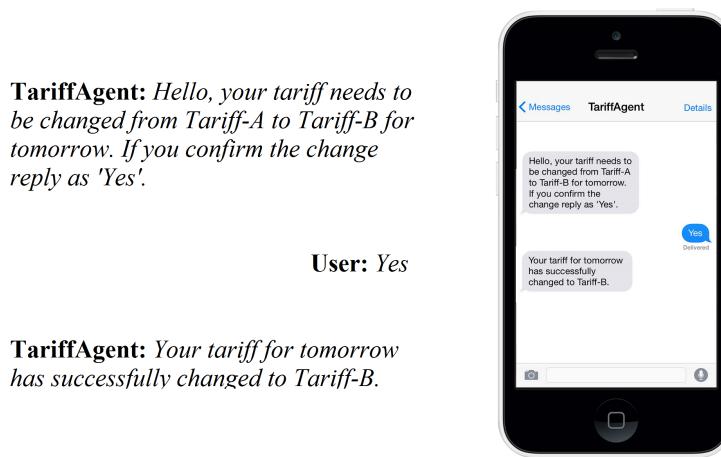


Figure 3.3: A text dialog under the semi-autonomous setting.

3.2.3.3 System Initiated Interactions via SMS

The system can send three different types of notifications via SMS: reports, suggestions and confirmations. Reports provide information on how much energy was consumed, how much the cost was, which tariff was selected, and how much was saved or lost (compared to the optimal or the worst tariff). Reports were sent every day to all users regardless of their setting. An example report is: “Hello, yesterday your tariff was Tariff-A, your consumption was 4.4 kWh and it cost you 0.69 pound. You saved 1.30 pound with Tariff-A.” The system sends suggestions to users who are on suggestion-only setting, when their tariff for the next day is predicted not to be optimal, for example: “Hello, your tariff needs to be changed from Tariff-A to Tariff-B for tomorrow. If you confirm the change please reply as ‘Yes’.”

Saving assumptions are only presented in the web UI for brevity, rather than in the SMS. Confirmation messages are sent only to users on semi-autonomous setting to inform them of an automatic tariff switch, such as: “Hello, I switched your tariff from Tariff-A to Tariff-B for tomorrow.”

3.2.4 Findings

We report findings from the semi-structured interviews through thematic analysis (Braun and Clarke, 2006). The interviews were documented through audio recording (later fully transcribed) and notes; analysis started by categorising the material at the sentence level through open codes. The codes were assigned to commonly recurring themes, significant events or references. The coding was performed by two researchers. Initially 88 open codes were used, which are later grouped in four broader categories that we discuss in what follows. We also present information on system usage based on automatic interaction logs.

Number of individual days	Alistia	Claudia	Ender	Greta	Ivan	Louisa	Maria	Mehmet	Omar	Sinha	TOTAL
Home page access	8	7	11	5	7	5	7	7	6	6	69
Details page access	5	5	11	3	6	4	4	4	4	4	50
Prediction adjustment	2	2	4	0	4	1	2	0	0	0	15
Manual tariff selection	2	3	2	0	4	1	1	0	1	0	14
Setting on semi-autonomy	0	0	5	2	1	13	7	0	0	0	28
SMS suggestion accepted	0	2	2	0	2	0	0	4	0	1	11
No response to SMS	1	3	2	2	1	0	0	0	1	0	10
Day of last interaction	14	14	14	14	14	12	14	14	14	14	

Figure 3.4: Overview of user activities. The rows listed under each user represent the total number of individual days that the user performed each activity. The last row shows the last day of user interaction.

3.2.4.1 Engagement

The summary of user activities throughout the study is displayed in Figure 3.4. Participants accessed the web interface on average every 2.2 days, with some participants accessing it almost daily and some as infrequently as once every 5 days. The home page was loaded more frequently, with 277 page loads over the course of the study, while the details page was loaded overall 210 times, still accounting for approximately 43% of the total page views (Figure 3.6).

The default autonomy level at the beginning of the trial for all participants was suggestion-only, where the system sends SMS suggestions about tariff switching but it does not automatically switch. This default option was chosen because it is the one that requires the most interaction from users, so we wanted to see whether they would change to a less demanding one over time. Half of the participants modified the settings to semi-autonomy, where the system automatically changes to the predicted best tariff and informs the user of the change via SMS. The semi-autonomous setting was kept for a maximum of 14 days, and 29 days in total across all the participants who used it (i.e., 5.8 days per participant on average). The remaining half kept using the default suggestion-only setting. No one selected the fully autonomous setting (where the system changes the tariff without informing the user). We report in Figure 3.5 for how the number of users keeping different autonomy levels varied over the course of the study.

In terms of tariff selection, two participants never received switching suggestions because they happened to be already on an optimal tariff. For the remaining eight participants who received tariff switching suggestions from the system (while on suggestion-only setting), five participants accepted them by replying ‘Yes’ via SMS at least once, the other three participants never accepted any tariff switching suggestion and stayed on the fixed tariff during the whole study. Overall, 15 suggestions were sent from the system to the users, with a 66% acceptance rate (Figure 3.7). All 5 participants who used semi-autonomous setting took advantage of the web UI to provide manual estimates of their electricity consumption prediction for the following day. In total this explicit input was provided 26 times during the study.

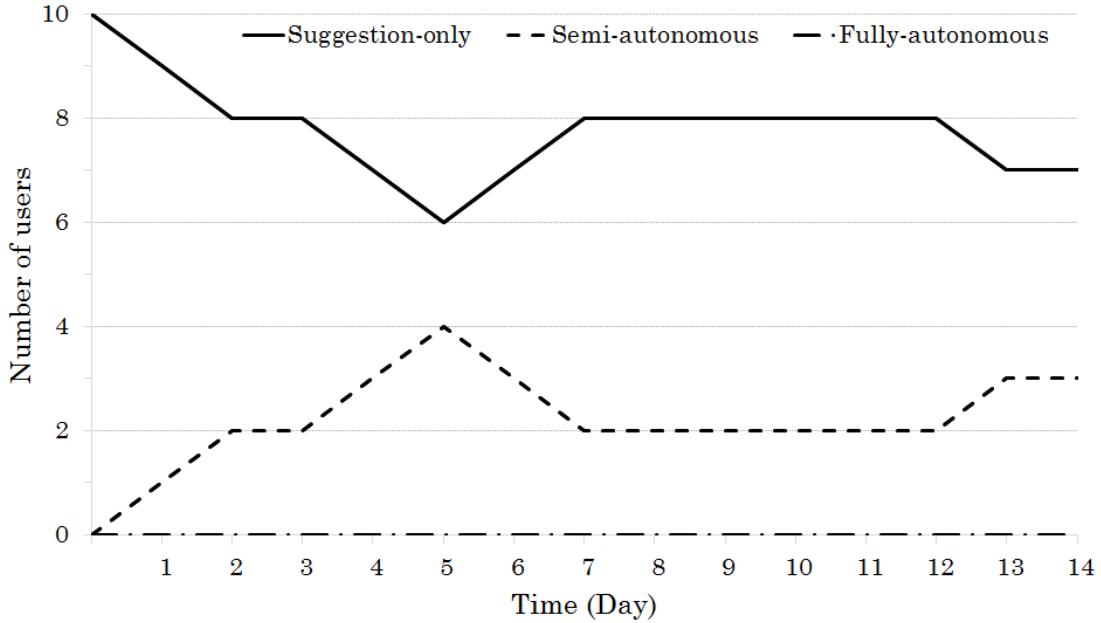


Figure 3.5: Autonomy level changes over 14 days.

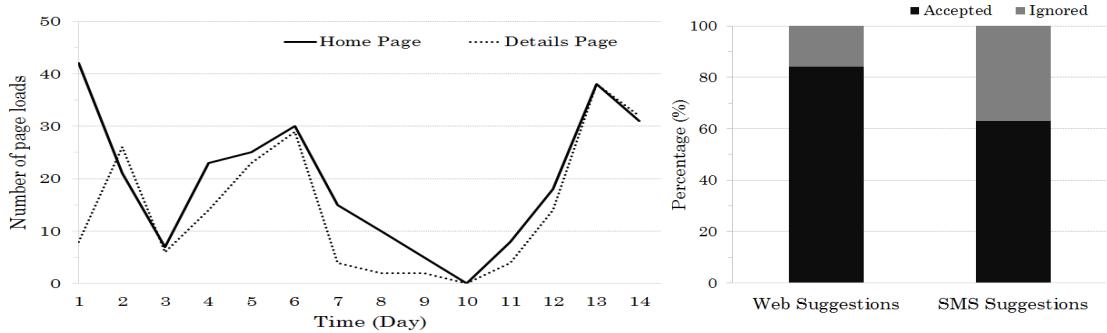


Figure 3.6: Page visits.

Figure 3.7: Acceptance of suggestions.

All participants stated that they checked their budget on the website with some regularity during the study, demonstrating that all participants cared about the reward. However, most of them reported that they more frequently kept track of their budget through SMS summary reports, as these reported the daily expenditure. Overall our participants spent £134.5 over the entire study, and they were rewarded a total of £165.5, corresponding to an average of £16.55 per participant (SD: £4.2).

All participants described the daily SMS notifications, which include the summary (tariff, cost and saving) of the previous day, as informative and motivating. Nobody complained about them being intrusive or too frequent. Moreover, the daily notifications were mostly explained as the preferred way to keep track of the system. For example, Maria said, “I like more to have an SMS than to login to the [web] page.” The amount of information provided in the SMS messages was reported to be concise and satisfactory. For instance, Ivan said, “It would be more confusing. I mean if I need more information to make a decision I would open the web site. [...] So, no, I think more information would be

more confusing.”, when we asked him if there was more information required in the messages. Some participants commented that they would opt for less frequent (e.g., weekly) summary messages if they had to use the system longer term:

Analyst - “What did you like the most and least about the system?”

Greta - “I think it is quite easy to understand, but after a while I would probably want the text messages to be less frequent. I would probably want to get updates every sort of a week maybe.”

3.2.4.2 Perception

Despite the limited duration of the study, the system appeared to be simple enough for some participants to develop sophisticated explanations about how the system and tariffs work, which were very close to the actual implementation.

Analyst - “What do you think the system does for you?”

Alisa - “Given a forecast of certain level of consumption which I tell to the system, it will select the best tariff, once it has selected the best tariff if I agree with that after counting the wind speed so splitting the cost in two parts. The one that is covered by the wind energy, once this part is consumed it will cost me the rest of the amount. I do not know if I explained it well. So in my mind, I thought like the actual amount of energy I have consumed there will be one part, which will be covered by the wind energy, and the rest of the energy will be paid by me.”

However, some participants had difficulties in understanding how the cost was calculated in the system. They referred only to the effect of their electricity consumption, leaving out the availability of wind energy.

Claudia - “It tracks how much electricity we use during a day and then suggests whether I should keep that tariff for the next day, or change to a different tariff. I guess it is sort of trying to see a pattern, every day we use a similar amount. I imagine it goes on the day before more. I do not know because I do not see how it suggested moving [to another tariff] on some days. Why did it suggest it, because it [the consumption] is not like being huge dramatically different each day?”

Some participants described the variable tariffs as risky since their rates can vary depending on consumption and available wind energy, whereas they are always charged at the same rate with a fixed tariff. In other words, with variable tariffs, they could end up paying more than what they were expecting (for a fixed amount of consumption), because of the wind energy that they have no control over. This appeared to be the main reason why three of the participants stayed on a fixed tariff for the entire duration of the study. In the interviews some of the participants reported being aware of the opportunity that they could save money by switching to the variable tariffs, but preferred avoiding risk.

Differently, we observed that some participants selected the variable tariffs not only because of possible monetary savings but also because they perceived the variable tariffs as more environmentally friendly compared to the fixed tariff. For instance, Alisa said, “I am saving money and it implies using wind energy. Because using environmentally friendly resources as wind energy is always good.” when we asked her why she mostly opted for the variable tariffs.

All participants commented that the system was functioning well and they all declared that they perceived the system to be working in their favour. In addition to being beneficial for the end users, some participants suggested that the system also works in favour of the energy provider as well as of the environment. Participants mostly found the system easy to use. For instance, Greta told us how it was easy for her to manage the tariffs even though she did not have any experience before.

Greta - “In terms of the online website, it is quite clear for me and also it shows you the budget. I think the quite good thing about the system is that you are updated through messages and you can sort of control the usage of tariffs as a person who has very little experience of dealing with electricity.”

The colouring of the history data in the detail page was not successful. Most of the participants had difficulty in remembering their meaning.

3.2.4.3 Adoption

Participants commented that the system reduces the hassle of dealing with tariffs and saves time and money. Being busy with travelling or other daily tasks was reported as a reason to select the semi-autonomous setting.

Analyst - “Why did you select the second setting (semi-autonomous)?”

Louisa - “I am very busy. I did not want one more task on my mind and I do prefer the confirmation by text because this way I know and I feel in control actually.”

However, even though participants appreciated the system autonomy, they were still keen to be aware of any possible automatic change. Therefore, all of the participants who delegated the system to automatically switch their tariff preferred using the semi-autonomous setting, where confirmation is sent via SMS (as opposite to fully-autonomous setting where no confirmation is sent).

Analyst - “Why didn’t you ever use the third setting (fully autonomous)?”

Greta - “I want to basically know what is happening. I used the automatically select the best tariff and send me confirmation because if something changes I would like to know about it.”

It is also interesting to note that the benefit provided by the system is often described in terms of offering control, even though autonomy is one of the major features of Tariff Agent, which was acknowledged by participants. For example, Louisa said, “[it] gives me [an] opportunity to save money, really, because *I can* change the tariff and *I can* predict for next day” (our emphasis). Through this quote we can see that Luisa, who mostly used the system in automatic mode, refers to herself, rather than the system, being in charge of switching tariff and achieving the savings.

In addition to the tariff selection, energy consumption is what participants have control over (while they have no control over wind energy generation). So the risk related to a variable tariff was sometimes associated to plans of carrying out energy-intensive activities, such as laundry.

Analyst - “How did you decide which tariff is the best option for you?”

Ender - “Actually, if you are planning to wash your clothes, it does not make sense to take risk because you know you will consume much more energy than yesterday.”

The participants who switched tariffs (either manually or automatically through the system) reported that they developed strategies to cope with the uncertainty of the wind and consumption. These were mostly based on planning future activities, and then taking advantage of the option to input their consumption prediction to inform the tariff selection by the autonomous system.

Analyst - “How did you use the button called ‘save into agent’?”

Maria - “This weekend was obvious because I was not here. So I decided to change the consumption prediction to be much less. Also I changed it on Friday, because on Saturday I use more energy usually than during the weekdays, because it is the day that I put on the washing machine and I stay more hours at home.”

During these days Maria kept the system setting to automatically switch tariff, so this quote demonstrates that taking advantage of system autonomy is not at odds with staying in control.

In the interviews, we received comments on how sometimes it could be difficult to predict the next day’s consumption, especially in a shared house.

Analyst - “Do you think the system worked well as it was intended?”

Alisa - “It [the system] would be very much useful for a household in which you can really make reliable forecast because for instance two days ago here there was a party of washing machine use. So you cannot play with saving into system setting, because it depends on other people’s decisions.”

Here Alisa highlights the challenge of using automatic predictions, or even manual predictions, given that household activities can sometimes happen in an unplanned fashion. The party of washing machine use that she mentions in this case refers to a day where her flatmates ran several laundry loads.

Some participants reported that receiving tariff switch confirmations from the system also reminded them to monitor the consumption prediction made by the system. For instance, Maria told us, “The best setting is the second one. The system selects the tariff automatically and sends me the confirmation that is it. Also, sending me confirmation reminds me to check if I will have a different consumption.” In contrast, others told us they did not often need to monitor the predictions as their consumption was more or less constant, for example Louisa said: “I just left it on automatic. I cannot say I was using it everyday because my consumption was kind of same.” This quote suggests that, in this case, the system autonomy was accepted, perhaps, because the prediction was particularly easy.

3.2.4.4 Accountability

All participants stated that they do not trust energy companies. Having bad experience with energy bills or hearing about companies in the media appear to be the main reasons

of distrust. Moreover, the abstruseness of energy contracts and energy bills increase the distrust of consumers towards energy companies.

On the other hand, all participants liked Tariff Agent, appreciated being able to change their tariffs more flexibly and having more transparency in their consumption and cost, therefore, found the system more reliable than what utility companies provide now.

Louisa - “With predicted stuff from the companies I always feel cheated but this way it is clear and it looks more honest really. Honestly, because it really feels like it is something transparent and straightforward.”

It was interesting to note that most of the participants felt that Tariff Agent offers more control, although the system works with a certain level of autonomy. This feeling of control seemed to increase the user trust towards the system.

Correct suggestions or tariff changes made autonomously seem to intrinsically improve the trust towards the system. For instance, Ender reported how a correct suggestion encouraged him to trust the system more for his future decisions: “It suggested me fixed tariff and I did not want to choose that and I lost money. In that case I learnt that I should trust the system.” Likewise, Alisa’s comment also suggests how the system accuracy plays a significant role in terms of trust: “I received the messages and the system actually was selecting for me the proper option, this is for sure. My tariff was always the best one, I saw it and I was happy about it.”

However, Tariff Agent was deliberately designed not to select the correct tariff all the time in order to elicit users’ reactions to an autonomous system that can negatively affect their budget. It was interesting to note that participants always referred to the system with tolerance when we asked how they felt about incorrect suggestions and selections in the interviews.

3.2.5 Summary and Implications

In the first study, it was interesting to note how different participants reacted to different autonomy levels, and appreciated automated SMS reminders and recommendations. Participants’ reports of their study experiences suggested that they felt more in control, engaged well with the system, and they were broadly tolerant to the system’s autonomous operations. This was particularly interesting since tariffs were confusing (i.e., wind-based tariffs added external uncertainty) and autonomous systems are generally mistrusted. This first study showed us that some participants willingly delegated their tariff decision to an autonomous system, but that they were still keen to monitor its operations and to intervene in the system when they believe it is necessary for improving its decisions.

However, the study was limited to two weeks. Even though the findings reveal that this two weeks was sufficient for participants to experience the system, form opinions about it and develop strategies to integrate the autonomous tariff switching in their everyday practices; the question naturally emerged whether the findings would also hold in a longer deployment period. In particular, was the user “tolerance” of the system’s autonomy and at times inaccuracies due to the scale of the study? Would users give the system more autonomy if they are satisfied with the performance of the system, or withdraw it entirely if they are not, as time passes? In fact, some participants reported that they might opt for the fully autonomous setting or would like to receive SMS reports less frequently, if the study was longer.

One common method for grounding the implications of a field study is to run a new study by utilising what have learnt in the previous one ([Sas et al., 2014](#)). Therefore, we decided to run a new longer field study to explore the questions stated above, after implementing some design changes for improving the first prototype based on the implications of the first study.

3.3 The Second Study

The second field study aimed to explore users’ longer term interactions with a tariff switching system. To recruit participants a similar approach was followed as for the first field study, however with an aim to have a broader sample of the population. Overall, 12 participants (6 female) were recruited to cover a range of lifestyles, as detailed in [Table 3.3](#).

In the light of the first study we implemented some design changes, which we detail in the following section. The participants used the new version of Tariff Agent for a period

Table 3.3: Participants’ profiles for the second study.

PARTICIPANT	GENDER	AGE	OCCUPATION	OTHERS
Adam	Male	60s	Priest	None
Arthur	Male	60s	Retired	None
Chloe	Female	40s	Community Manager	None
Dionisia	Female	20s	Chemistry PhD Stud.	1 Adult
Evelyn	Female	60s	Retired	None
Gerard	Male	20s	Media Production Manager	1 Adult
Gonca	Female	20s	Social Policy MSc Stud.	None
Hiroko	Female	30s	Housewife	1 Adult, 2 Children
Lucy	Female	20s	Estate Manager	1 Adult
Peter	Male	30s	Unemployed	1 Child
Stewart	Male	30s	Software Consultant	1 Adult
Turan	Male	30s	Chemistry PhD Stud.	1 Adult

The screenshot shows the TariffAgent home page with the following sections:

- Tariff:** Your current tariff is **Tariff-A**. Tomorrow's tariff is **Tariff-A**. You can change it before 9pm today.
- Predictions for Tomorrow:**
 - 4.1 kWh - A lot more than yesterday
 - 3.2 kWh - More than yesterday
- Estimated Consumption:** 2.7 kWh - **Same as yesterday**
 - 1.9 kWh - Less than yesterday
 - 1.4 kWh - A lot less than yesterday
- Tell the Agent** button
- Setting:**
 - Send me an SMS when tariff change is suggested
 - Automatically select the best tariff and inform me
 - Automatically select the best tariff without informing me
- Reports:** I want to receive a report every:
 - day
 - 3 days
 - 5 days
 - week

Reports are sent regardless of your selected setting and they include brief information about your consumption, cost and selected tariff.
- Budget:** **Available:** £59.70 **Spent:** £20.30
- Notes:** Please [click here](#) to see your account details.

Figure 3.8: New home page.

of 6 weeks (42 days). At the beginning of the study, all participants were allocated a budget of £80, and their daily consumption cost was reduced from this budget over the period of the trial.

3.3.1 Implemented Changes

During the interviews of the first study, we observed that some participants perceived wind-based tariffs as more environmentally friendly. This perception influenced their reactions to the tariff suggestions that were automatically sent by the system, and led them to ignore the suggestion and stick to the same tariff that was not always the cheapest. Therefore, we removed the emphasis on renewable energy and external uncertainty to more distinctly focus on issues of users' orientation to smart systems. Similarly to real-world tariffs, the tariffs we use in the second field study are defined in terms of: a standing charge that is a fixed amount charged daily for service cost; and a unit rate that is the price of electricity per kWh. In particular, each tariff represents the best value for a particular consumption range so that it is not easy to decide which tariff is the cheapest as it may change every day unless the user is able to accurately predict her own consumption.

In the first study, the system was designed to send a daily report to all users regardless of their settings and the users were not able to deactivate or reduce the frequency of

TariffAgent							
Account Book Details							
Date	Predicted Consumption (kWh)	Actual Consumption (kWh)	Agent Suggestion	Selected Tariff	Budget (€)	Cost (€)	Saved/Lost (€)
31-Aug	3.4	4.3	Tariff-A	Tariff-A	65.71	0.68	1.313
30-Aug	3.4	5.1	Tariff-A	Tariff-A	66.47	0.76	-0.002
29-Aug	3.6	3.4	Tariff-A	Tariff-A	67.06	0.59	1.394
28-Aug	2.8	3.4	Tariff-A	Tariff-A	67.65	0.59	1.394
27-Aug	4.5	3.6	Tariff-A	Tariff-A	68.26	0.61	1.376
26-Aug	2.9	2.8	Tariff-A	Tariff-A	68.79	0.53	1.448
25-Aug	5.3	4.5	Tariff-B	Tariff-B	69.50	0.71	-0.010
24-Aug	3.0	2.9	Tariff-A	Tariff-A	70.04	0.54	1.439
23-Aug	3.3	5.3	Tariff-A	Tariff-A	70.82	0.78	-0.006
22-Aug	4.3	3.0	Tariff-A	Tariff-A	71.37	0.55	1.430

█: Incorrect consumption prediction that caused incorrect tariff suggestion
█: Incorrect consumption prediction that did not cause incorrect tariff suggestion
█: Correct consumption prediction that lead to correct tariff suggestion

◀ Previous ▶ Next

Figure 3.9: New details page.

these reports. In the interviews some users commented that they could opt for a less frequent summary messages, for example weekly, if the study was longer. Therefore, we decided to alter the design and enable users to change the frequency of reports. To do so, we added another component to the home page, called Reports (see Figure 3.8). This component enables users to decide how often they want to receive an SMS report, where the option every day is initially selected by default. The other options are: every 3 days, every 5 days and every week.

Furthermore, in the home page we added explanatory information tips to help users to understand and remember the basic underpinnings of the system, for example what the functionality of ‘Tell the Agent’ button is, how the estimated cost is calculated, and what information the SMS reports include.

From the details page (Figure 3.9), we removed predicted and actual wind energy values in accordance with the changes in tariff specifications. We added a Budget column, showing the balance left in their account for each of the past days. Also at the bottom of the page we inserted colour-code reminders that explain how the table values are coloured. For instance, a green colour represents a correct consumption prediction which resulted in a correct tariff suggestion (but not necessarily a correct tariff selection since the user might not have selected the suggested one).

3.3.2 Findings

Similar to the first study, we used thematic analysis (Braun and Clarke, 2006) for the semi-structured exit interviews. The same two researchers performed the coding. Initially 76 open codes were used, which are later grouped in four broader categories that

Number of individual days	Adam	Arthur	Chloe	Dionisia	Evelyn	Gerard	Gonca	Hiroyo	Lucy	Peter	Steward	Turan	TOTAL
Home page access	11	18	6	9	19	26	16	12	4	4	6	4	135
Details page access	9	1	0	3	8	21	9	4	1	3	6	2	67
Prediction adjustment	3	16	0	7	7	25	3	6	1	4	1	3	76
Manual tariff selection	3	17	2	7	9	1	1	2	1	4	5	4	56
Setting on semi-autonomy	39	0	0	0	0	40	30	0	0	0	13	0	122
SMS suggestion accepted	0	0	5	2	3	0	1	0	23	5	2	11	52
No response to SMS	2	0	10	13	21	0	2	9	2	4	14	10	87
Day of last interaction	39	37	41	38	41	38	31	24	42	41	29	42	

Figure 3.10: Overview of user activities. The rows listed under each user represent the total number of individual days that the user performed each activity. The last row shows the last day of user interaction.

we discuss in what follows. We also present information on system usage based on automatic interaction logs.

3.3.2.1 Engagement

The summary of user activities throughout the study is displayed in Figure 3.10. Participants accessed the home page more frequently, with 420 page loads over the course of the study, while the details page was loaded overall 184 times, still accounting for approximately 30% of the total page views. Figure 3.12 shows that the interactions of users with the web site drastically dropped after the first week, yet were maintained at a significant level by each user until the end of the study. Everyone either accessed the web interface or replied to SMS suggestions with some regularity that is on average at least once every 2.5 days (SD: 1.3 days).

Similar to the first study, the default autonomy level at the beginning of the trial was suggestion-only for all participants. Four of the participants modified the autonomy level to the semi-autonomous option, where the system automatically changes to the predicted best tariff and informs the user of the change via SMS. The remaining eight users kept using the default autonomy level. No one selected full autonomy (where the system changes the tariff without informing the user). We report in Figure 3.11 for how the number of users of each autonomy level varied over the course of the study.

In terms of tariff selection, the system’s predictions of the optimal tariff were correct on average 65% of the time (SD:12.3%); for reference users’ actual selections throughout the study (a combination of automatic and manual) corresponded to the optimal tariff 61% of the time (SD:15.4%). Overall 10 participants received SMS tariff switching suggestions from the system (while on the suggestion-only setting). The other two participants did not receive any tariff change suggestions, as one was not using his mobile phone and the other one switched to an automatic tariff change setting after the first few days. Eight participants accepted the suggestions by replying ‘Yes’ via SMS at least once, the other two participants never accepted any tariff switching suggestions but changed their tariff manually on the website at least once. In total 139 SMS suggestions were

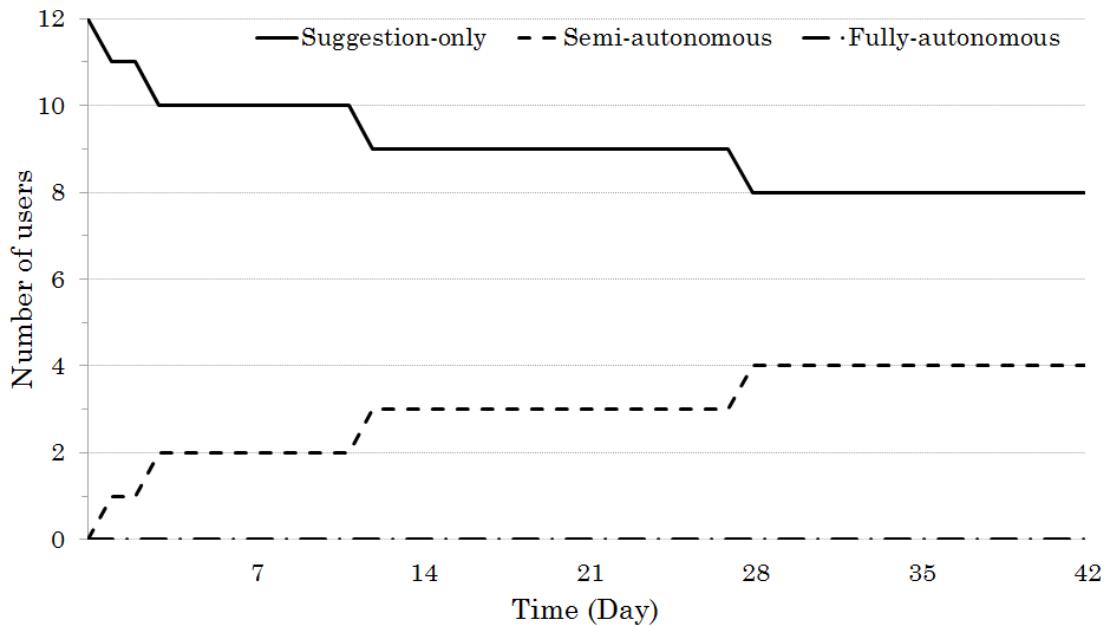


Figure 3.11: Autonomy level changes over 42 days.

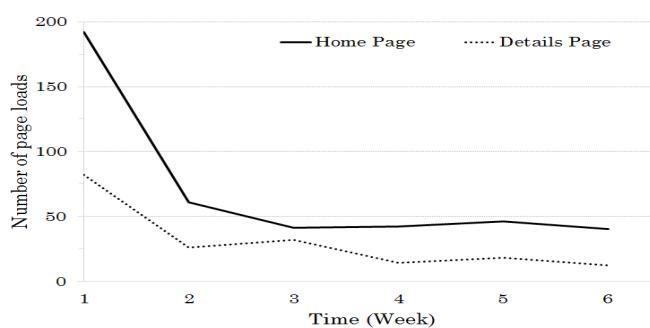


Figure 3.12: Page visits.

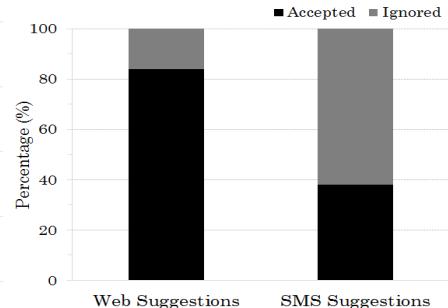


Figure 3.13: Acceptance of suggestions.

sent from the system to the users, with a 38% acceptance rate. Additionally, 59 times users changed their tariffs manually from the website and in 49 cases they accepted the system's suggestion (see Figure 3.13). In these cases an SMS was not sent because the suggestion was accepted early in the day, via the web UI.

All users except one took advantage of the web UI to provide manual estimates of their electricity consumption prediction for the following day at least once. In total this explicit input was provided 110 times during the study, and resulted in 85 correct tariff selections. The one person who did not use the input feature turned out to have a very regular consumption profile throughout the entire duration of the study.

During the interviews, participants stated that they checked their budget regularly from both the website and the SMS reports, which shows that they cared about the reward. Overall our participants spent £440 over the entire study, and were rewarded a total of £520, corresponding to an average of £43.30 per participant (SD: £14.67).

All participants commented that the system was functioning well and that they found it easy to use. The following quote is representative of the sort of reactions we recorded: “I think it was easy to use. It was really simple, a few buttons to click, and if it needs to change the tariff, hmm getting the text alerts I quite liked.” (Peter). The quote demonstrates also the participants’ general appreciation for the SMS notifications. We were often told that they serve as reminders and as an easier way to access the system, compared to visiting the website.

Similar to the first study, participants described the daily SMS reports, which include the summary (tariff, cost and saving) of the previous day as informative and useful. Nobody found them intrusive or too frequent. In fact, interestingly all participants kept the frequency of reports at the daily option, which was default. Three users lamented that they found the saving/loss information in the SMS reports confusing, because it was not clear to them what their expenditure was compared against.

The value of SMS reports and notifications was emphasised by one participants who did not have a mobile phone during the study: “I do not have a mobile phone, so the text message was coming through on my landline [through a text-to-speech service]. It could have been clearer by an email; I look at emails every day so there is no problem there. So for me, that would have been an improvement on how it worked.” (Arthur). The suggestion to use email as an alternative to SMS indicates that Arthur felt the lack of notifications and reminders, even though he was a very active user of the web UI (he manually changed tariff on 17 days and adjusted the consumption prediction on 16 days).

3.3.2.2 Perception

Some of the interview questions were aimed at assessing participants’ perception of the system. Most participants appeared to hold a mental model that mirrors quite closely the actual design and implementation of Tariff Agent, except for the consumption prediction as detailed below. For example, Gerard eloquently described: “It looks at what usage you think you are going to be using in the next day, which is something you either tell it or it just guesses itself, and then it looks at various tariffs based on the standing costs and the per unit costs, and works out which one is going to be the cheapest.”, or more concisely Hiroko told us: “It decides the tariff and recommends me the tariff.”

However, for some other participants the initial response to the question of “what the system does” instead highlighted energy monitoring and awareness. Peter told us: “Well it just helps me to monitor my energy usage, it would help me definitely bring my bills down. That’s what I was hoping it would do, anyway. That’s what it seems to do. It definitely helped me be more careful with my energy usage.”

Indeed, some participants reported that such awareness led them to change their energy consumption habits, for example switching off devices that used to be left on or in stand-by.

Analyst - "Have you changed anything related to your energy consumption?"

Lucy - "I used to leave the telly on for my dog during the day but I do not do that anymore. He sits in silence now. It is just as I said probably made me aware how much electric I actually use whereas before I just used to put 10 pound a week electric and just done with it. Now I am thinking actually seeing where things going. So it has made me changing definitely."

Through the interviews we could also notice that some participants had a quite detailed model of how energy tariffs work.

Analyst - "How do you think your daily energy cost is calculated?"

Chloe - "I reasonably understand that yeah, higher standing charges for higher tariffs, but because my unit rate is low then I am better off having either Tariff-A or Tariff-B with the lowest standing charge."

Our scenario, where the tariff can be switched everyday, was considered by some as a key feature that is part of Tariff Agent.

Analyst - "What did you like the most about the system?"

Arthur - "I did like the different choices. It was quite interesting having some, quite variation on unit rate and standing charge can compare how they would produce the actual cost."

Towards the end of the interview, participants were asked who they thought would own the system, if it was real. The answers were quite varied, but most participants seemed to have quite a strong view on the matter. Many suggested that they would personally prefer the government or an independent organisation to provide the system. For example, "Ideally, somebody who has got no vested interest in what tariff it is." (Evelyn). Or, "I could not image this implemented by anybody other than power companies realistically. But I personally prefer government." (Gerard)

3.3.2.3 Adoption

The interviews revealed a range of attitudes and orientations towards the autonomy of the system, which showed similarities to the ones that we observed in the first study. Being busy with travelling or other daily tasks, and feeling lazy were reported as reasons to select the semi-autonomous setting, for example, “I was logging in everyday and just looking at it and leaving it and to be honest towards the middle of it I kind of got a bit more lazy.” (Stewart)

Two of the participants who used the semi-autonomous setting reported to have already been exposed to real-world systems that recommend changes of tariff via email, “I actually get an email from them if a cheaper tariff based on my last month usage is available. So if they find a tariff that would save me more than £50 a year I get an email.” (Gerard)

Another participant highlighted the options offered by Tariff Agent in terms of automation against manual intervention, and appreciated being able to shift between these multiple options.

Analyst - “Can you describe how you used the system?”

Adam - “I am very, very busy here within the parish. So I was fascinated by the way in which you have a multiple choice [setting]. You could either how to say look at it, examine each aspect and dimension of this program, or you could choose as I did after a day or so number 2 [semi-autonomy], and it would choose it automatically, and I found that great help as well.”

Along the same lines, participants generally considered the opportunity of providing input to the system as an advantage. This aspect was generally related to the fact that it would be impossible for an automatic system to predict correctly in case of exceptions in their daily routines, “I think the whole idea of the fact that you can have an input is the reason why I like it.” (Lucy), or “That is the point of the system, isn’t it really, that you monitor it and you go no hang on that is not quite right. So ultimately you have to be responsible for it. The control is still with user ultimately.” (Peter)

However, the input that users need to provide to the system was perceived as effort that needs to be put into using it. Some participants perceived it as light, “It does require you to do a bit of work, but I do not think it requires you to do an awful lot of work.” (Gerard), while others considered it hard, “OK but then that involves a lot more input from me to do the prediction for what tomorrow might be.” (Chloe). In some cases, the increased performance provided by the manual input was not considered to be worth the effort, compared to the performance that Tariff Agent would achieve on its own.

“Every other day it gets it right, every other day it gets it wrong and then to be honest, I think I would probably make the same mistakes so and I am not going to login in every single day and change it because it is more effort than it is worth I guess in that case.” (Stewart)

Some participants reported that they kept the system in suggestion-only setting because they generally like to be in control. For example Dionisia told us: “I think it [suggestion-only] suits me. I like to be in control of what I spent, I guess.” Similarly, Peter said: “I think I quite liked the idea of still being in control of it. So still being, yeah so it was still my choice rather than letting you to decide.” Disagreement with the choice of the system was also mentioned as a reason not to relinquish control: “I would never use automatically select ever. [Laughing] ’cause I may not agree.” (Evelyn)

As described in Section 3.2.2, Tariff Agent uses the previous day’s consumption to predict for the next day, and estimate the best tariff. Users have the option to adjust this prediction through the web UI. Indeed, most participants took advantage of this option and reported that they predicted the next day’s consumption mostly based on planning future activities such as doing laundry, being away from home, or hosting guests. Changes in personal plans were reported as an obstacle for providing input to the system: “I found it impossible to predict tomorrow because I might make plans to use washing machine or a cooker and then I didn’t do that.” (Evelyn)

On the other hand, another participant told us how she changed her plans to match the tariff selected by the system.

Gonca - “Sometimes I could not change the tariff when the system showed me like tomorrow my tariff will be changed to B, which was more than I consume normally, then I knew that ok today I will not do the laundry, for example, and I just waited to do the next day.”

In this case Gonca reported to have seen the system notification too late, after the 9pm deadline for tariff selection, so she changed her plans for the next day to follow the system schedule.

Some participants suggested that contextual factors specific to their circumstances influenced the accuracy of predictions. For example, Adam said, “This house unfortunately has so many visitors coming through that when the choice was made based on yesterday, I’d get a lot of unexpected visitors coming and really [laughing] test the patience of this program.” Similarly, Evelyn indicated that predicting could be even harder for low profile users since any appliance usage might cause peaks in the consumption: “In my case, my consumption is going to vary greatly from day to day, with a family that might be different.”

A number of participants reported that they were confused with the system being one day behind and one day ahead. For example, Peter said, “Sometimes it was difficult. I get confused sometimes thinking about yesterday in relation to what energy I’d use tomorrow, but obviously that’s the only way you can do it, because yesterday you had the all-day consumption, so that was just me trying to get my head around it.”

We also recorded a number of comments related to the mechanics of the prediction system. While some users correctly understood the system to predict tomorrow’s consumption simply to be the same as yesterday’s consumption, a number of other users expected the system’s prediction mechanism to be more sophisticated than it actually was. For instance, Arthur told us that the system was learning his consumption: “The pattern that has been build up over the previous days or weeks. Of course for the longer period, then hopefully it would get to know how my energy usage comes out in regular basis.” The same participant also commented that the system may help him to improve his predictions over time: “It would also, perhaps, get me to estimate better what my consumption is going to be so I think I would learn as I went along more how to estimate the consumption.”

In the same vein, talking about how the system predicts her consumption for the following day, Lucy explained: “Probably what I have used on the same day, or I suppose as it goes on longer, you get more data so you got more of an idea of what we do use and what we don’t use. So I think probably as time got on it was more accurate than it was initially ’cause you collect more data.” Others were uncertain about the matter: “I don’t know. I was wondering how smart it was, whether it looks at previous weeks, days of the week.” (Gerard)

3.3.2.4 Responsibility

Similar to the first study, all participants commented that they perceived the system as helping them save money, through mostly correct suggestions, and most of them stated that they trust the system’s tariff decisions. However, the participants of the second study were more aware of the system’s imperfection, most probably due to a longer duration of the study.

Interestingly, when we asked about experienced mistakes in tariff suggestions or selections, they mostly considered it to be their own responsibility, “It is my mistake because I have not informed the system of something that I knew would change what consumption I used.” (Lucy), or, “it is connected to the individuals and requires individual responsibility to make the system in the savings work, it sort of gives responsibility back to the energy user.” (Chloe) In an even more drastic way Gerard said, “I think it was always my fault. So I just felt annoyed with myself and I started trusting the machine

more than I trusted myself, I was like oh the machine knows what it is doing, just leave it alone.”

Talking about mistakes the system may commit, Arthur highlighted the importance of receiving immediate feedback in order to maintain trust, “I would not worry as long as I got the message straight away to say our suggestion yesterday was erroneous.”

3.4 Discussion

In our two field studies we exposed two diverse groups of participants to an envisioned future energy scenario, in which autonomous systems embedded in households have the ability to switch the energy tariff based on the offered rates and a prediction of the user’s consumption. However, to reveal meaningful results from such studies it is essential that participants feel and engage with the scenario. Therefore, to render the scenario tangible for our participants, we combined a financial experimental reward with actual energy consumption data measured in participants’ homes. Both user interaction logs and the accounts offered in the interviews about system usage suggest that the study design was successful in making the scenario visible and tangible, echoing recent findings from the literature ([Costanza et al., 2014](#)).

3.4.1 Interaction with Autonomy

One aim of our research was to explore how users would perceive an autonomous system affecting them financially and possibly intruding into their daily routines. The results of the first and the second study suggest that participants kept a strong feeling of control over the autonomous system and they appear to understand and appreciate the autonomous nature of Tariff Agent, which proactively sends users suggestions about switching energy tariffs. In particular, this is well demonstrated by the general and prolonged engagement that most participants of the both studies had in switching tariffs and in providing input to the system to improve its consumption forecast performance, and hence the quality of the suggestions. Additionally, none of the participants reported that they found the SMS tariff switching prompts bothersome, supporting the findings of other recent work ([Bourgeois et al., 2014](#)). Also, as suggested by some participants of the first study, we implemented the second prototype to enable users to change the frequency of summary reports sent daily by the system. However, no participants changed the frequency of the reports and all kept it at daily option, since they found the SMS reports informative and useful.

Our other aim was to investigate the users’ autonomy and interaction preferences over both short and longer terms. In the second field study the engagement with autonomy

delegation was less generalised compared to the first study. While half of the participants used the semi-autonomous option in the first study, only one third of participants delegated responsibility to the system to automatically switch their tariff in the second study. This finding indicates a key design implication for autonomous energy systems. It is important to offer autonomy delegation as an option, as some users took advantage of it. This is also supported by the diversity of responses we collected around participants' orientation to autonomy. In other words, the decision of balancing control and autonomy should be left to users where appropriate to cater for individual differences in what levels of system autonomy people are comfortable with. In so doing, users may continue to make use of relevant parts of the system autonomy, rather than taking over full control or abandoning the whole technology (Yang et al., 2014). Our results suggest that flexible autonomy shows promise for sustaining users' engagement with an autonomous system, despite its occasional mistakes.

For both studies, the interviews revealed that all participants maintained a strong feeling of control over the system. It is worth emphasizing that the duration of the trial and the presence or absence of renewable energy do not seem to influence this finding. This result becomes more interesting when we consider that existing literature suggested that autonomous technologies leave users feeling out of control (Barkhuus and Dey, 2003), and that in our application even with the lowest autonomous level (i.e., suggestion-only) the system was continuously monitoring users' consumption and automatically suggesting tariff changes. Such a feeling of control seems to be boosted by the awareness that the autonomy level can be changed at any time, that input about their own consumption predictions can be provided to the system to help it make better decisions, and that they are able to easily monitor the performance of the system. This is particularly shown by the participants, with most feeling that incorrect tariff decisions would be largely their responsibility. Recent work Stout et al. (2014) suggested that delegation of autonomy is perceived by users as a means to shift blame from themselves to autonomous system for undesirable outcomes. In contrast, our results show that the ability to delegate autonomy in a *flexible* manner ultimately leads users to feel in control and therefore responsible.

Such a result may be, in part, due to domestic energy consumption being perceived as a very personal matter, over which an automatic system has limited insight. Nevertheless, we believe this is an important finding, which suggests that flexible autonomy may lead people to share the blame with the system and form more collaborative and long term relationship with it, rather than only blaming the system for incorrect actions and stop using it. This is also a promising finding for previous HCI studies (Strengers, 2011; Darby, 2006) that address the challenge of maintaining user interest and interaction with the eco-feedback systems. It looks like flexible autonomy could be useful to sustain user interaction and may help them to adapt new behaviours in longer run.

3.4.2 Orientation to Smart Systems

As detailed in our system description, we decided to implement a very simple consumption prediction strategy in Tariff Agent because we expected it would make it easier for users to interpret the functionality of the system, while offering comparable accuracy to more complex alternatives. Namely, the system was simply using the previous day's consumption as a prediction for the next day. Although this simple strategy is even indicated on the system's home page, in both studies some participants perceived the system's predictions to be smarter and more complex than they are in reality. They suggested that the system applies more advanced pattern recognition (e.g., days of the week) and even that it learned their usage habits, becoming more accurate over the duration of the study. Once again, this finding emerged irrespective of the duration of the study and the presence of wind-generated energy. This result resonates with the findings of (Yang et al., 2014) around their study of a commercial smart thermostat: their participants overestimated the abilities of the product, for example in terms of learning when they are at home and when away, leading to some resource wastage.

While in the case of the smart thermostat it may be explained through the explicit marketing of the product as "learning over time", we did not promote Tariff Agent as having such features. This expectation, then, may be due to an ever-increasing exposure of the general public to advanced machine learning computer systems. Such exposure is sometimes direct (e.g., Internet search engines), or through media reports. An alternative explanation may be related to our innate talent to learn over time. Perhaps participants expected Tariff Agent to learn like a person would. Even though we carefully avoided any anthropomorphic feature, to avoid any emotional biases, the system, even by its name, is explicitly referred to as an "agent" which may create such expectations. We believe these results highlight an important implication for future research in interaction with "smart" systems: to try and discover the source of people's learning expectations.

In the second study, participants comments around who would, or should, own the system if it was real further suggest that Tariff Agent seems to be perceived as more opaque than we had hoped for. Often the preference was for the ownership and operation to lay with the government or with an independent body. We interpret this as an indication that participants are afraid that the performance of a system dealing with energy consumption, tariffs and billing, may be biased to favour the profits of an energy provider. The system is not perceived as a simple and neutral data processor, nobody answered to the ownership question saying that it does not matter. This finding resonates with those reported by (Rodden et al., 2013), who exposed users to two different video scenarios where an autonomous system mediated energy and tariffs, and found that when the system was presented as being installed and owned by the energy supplier, reactions were much more critical than when it was presented as being installed and owned by the householders themselves. The perceived opacity is at odds with the fact that, at the

same time, our participants felt largely responsible for any losses incurred by incorrect tariff selections, rather than blaming the autonomous system. This apparent contrast can perhaps be explained in light of the fact that in our studies the system was clearly presented as developed and operated by a university, a type of organisation generally recognised as trustworthy.

3.4.3 Design Implications

In what follows, we list design suggestions based on our in the wild evaluations and qualitative and quantitative analysis for specifically autonomous domestic energy systems, where the systems deal with financially sensitive tasks and may disrupt users' daily activities. However, we believe that these design guidelines can be also exploited for developing applications in other domains that involve human-agent interaction.

- *Provide an easy way for users to receive updates about the status and operation of the autonomous system, and allow users to change the frequency of the updates.*

This first suggestion is based on the observation that in our field trials participants were very keen on keeping track of the system's operations. None of them disabled the SMS notifications; instead everyone reported that they found them useful. Moreover, the web access logs illustrate that participants visited the detailed information page quite frequently; it received about 43% of total page views in study 1, and 30% in study 2, illustrating a desire to monitor in detail what the system is doing. Although none of our participants reduced the frequency of daily summary SMS messages in the second study, providing such ability to users seems useful as some participants suggested so in our initial study.

- *Enable users to instruct the autonomous system by offering them opportunities to declare their plans and integrate these plans into the system's operation.*

This second suggestion builds on the perception of feeling in control that our participants reported in the two trials. This feeling of control was related to them inputting into the system their predicted consumption for the following day, but also to some adjustments of the participants' schedule. Indeed in some cases this action took the form of almost "booking" their activities (e.g., laundry) into the system. Moreover, being able to correct the system's prediction seemed to improve the intelligibility of the system's decision-making as this created a sandbox (Mennicken et al., 2014) for users where they could try different consumption inputs to see how the system accordingly changed the suggested or selected tariff.

- *Leave the system open to transfer of control by allowing users to adjust the system's level of autonomy (i.e., when to release or retain autonomy).*

This last suggestion is based on the fact that in both field studies some participants used the autonomous system setting of Tariff Agent (semi-autonomy), while other

users kept the manual confirmation (suggestion-only). Of those who took advantage of automation in the first trial, two users reverted to suggestion-only mode after some time. It should be emphasized that this suggestion-only option does not correspond to simply turning the system autonomy off completely, indeed the system still continuously monitors consumption and it autonomously offers suggestions for when to change tariff. However, the user needs to explicitly accept such suggestions, before they are turned into practice. Results from both studies show that individuals might have different preferences for different autonomy levels, and this preference might change over time.

These design suggestions extend and enhance the existing trends in mixed-initiative systems. These put the onus of requesting human control or input on the system (Horvitz, 1999; Scerri et al., 2003). Mostly it is part of the system’s functionality to decide when to attempt and transfer control to users. In contrast, we suggest that a more suitable approach is to enable the user to enact control or to adjust autonomy by default. We believe that systems involving humans and agents (so-called human-agent collectives (Jennings et al., 2014)) should enable human users to provide input to autonomous systems at any time to improve their operation. In turn, this makes legibility of system state an essential requirement for the design of mixed-initiative systems. In order words, we suggest that systems involving humans and agents should be left open enough that users can decide when to intervene. It should not be necessary to express this operation as “removing” or “diminishing” agency from the system, indeed in our system the optional user input provides more information for the system to help them to save money.

3.5 Summary

In this chapter, we have presented two field studies that exposed participants to a prototyped future energy scenario. Our scenario simulates a situation where households can switch electricity tariff on a daily basis, to try and best match their consumption level. This scenario enabled us to study users’ interactions with Tariff Agent, an interactive autonomous system designed to help in managing energy costs, which offers flexible autonomy and detailed information about its operation. The studies were made possible by a medium-fidelity prototyping approach combining off-the-shelf Internet-connected sensors with Web technology, and with financial experimental rewards.

Our field studies enabled participants to experience an autonomous energy system in their everyday lives, form opinions about it and develop strategies to integrate its autonomous operation in their everyday practices. The results of our field studies demonstrate that users are happy for the help Tariff Agent provides to them to deal with the complexity of variable tariffs, and at least in part, ready to use systems like Tariff

Agent to manage their energy tariffs. Our results suggest that people are willing to delegate some control to the agents but not fully. Hence, it is important to stress that system designs need to strike a nuanced balance between providing the user with means to monitor system performance and take control when they consider it necessary. Our results highlight opportunities and show promising directions to design mixed-initiative autonomous energy systems. In particular, based on the results, we have provided design guidelines for developing future intelligent energy system to make these system useful and acceptable to users in their everyday lives.

These design guidelines presented in Chapter 3 are based on the use of a simple agent whose operations are fairly intelligible. However, we have not studied the further question: how to design interactions of a learning agent to meet the requirements stated in Chapter 1. In particular, making the actions of a learning agent intelligible is challenging due to the fact that the agent will be learning how to react to the user's input over time (unlike Tariff Agent where users' consumption input had immediate impact on the agent's tariff decision). Therefore, the user's input might not directly manipulate the agent's actions as that single input could be treated as noise depending on the prior knowledge, which may eventually lead the user to think that the agent is malfunctioning. To address this challenge, in the next chapter, we represent a study of a learning agent called Smart Thermo, which helps users to manage their home heating control given real-time prices.

Chapter 4

Smart Thermo

In this chapter, we aim to explore the design space of a smart thermostat that helps domestic users react to real-time energy pricing, by autonomously adjusting the indoor temperature setpoint. More specifically, we aim to address the research questions: what would be people's feelings and expectations towards a smart thermostat that controls their home heating based on real-time prices, whether different user interface designs of the thermostat result in different user understanding and reactions, and how people would adopt and interact with the thermostat in their everyday lives.

To this end, we designed, implemented and deployed three different designs of a smart thermostat: a rule-based thermostat by which participants manually specify how to respond price changes, and two learning-based thermostats that apply a machine learning algorithm to identify households' temperature preferences over different prices. In order to evaluate these designs and to observe how users would react to such *future technology*, at a point in time where energy real-time prices are not yet widely implemented, we conducted a field study based on the same methodology presented in Chapter 3, where we exposed our participants to a future scenario through a combination of financial experimental reward and sensors installed in their homes. As our smart thermostat responded to the varying prices on households' behalf by adjusting the home's temperature, it caused a real impact on the comfort of its owners.¹

4.1 The Study

We conducted a study with 30 UK households (see Table 4.1) over a period of four weeks during February-March 2015. To recruit participants we distributed approximately 3000 study invitation letters around the city. We recruited households who had a broadband

¹Ethics approval reference for the study is ERGO-13417.

Table 4.1: Participants' profiles.

	Thermostat	Age	Occupation	Others
P1	Indirect L.	40	PhD Student	1 Child
P2	Manual	63	Maintenance Eng.	1 Adult
P3	Direct L.	37	Antiques Dealer	1 Ad., 2 Ch.
P4	Direct L.	43	PhD Student	1 Ad., 2 Ch.
P5	Manual	50	Estate Mng.	1 Ad., 1 Ch.
P6	Manual	36	Nanny	1 Ad., 1 Ch.
P7	Indirect L.	76	Retired	1 Adult
P8	Indirect L.	32	Radiographer	2 Ad., 1 Ch.
P9	Direct L.	44	Teacher	1 Ad., 2 Ch.
P10	Direct L.	62	Retired	2 Adults
P11	Manual	44	Teacher	1 Ad., 2 Ch.
P12	Indirect L.	50	Office Mng.	1 Adult
P13	Manual	40	Education	2 Ad., 1 Ch.
P14	Indirect L.	60	Lecturer	1 Adult
P15	Manual	53	Photographer	2 Ad., 1 Ch.
P16	Direct L.	60	Self Employed	1 Adult
P17	Indirect L.	58	Charity Mng.	1 Adult
P18	Manual	40	Accountant	1 Ad., 4 Ch.
P19	Direct L.	71	Retired	1 Ad.
P20	Indirect L.	56	Database Admin.	2 Adults
P21	Indirect L.	26	Contract Mng.	1 Ad., 1 Ch.
P22	Direct L.	22	Student	1 Adult
P23	Indirect L.	28	Sport Mng.	2 Adults
P24	Direct L.	69	Retired	1 Adult
P25	Direct L.	49	Gas Engineer	1 Adult
P26	Manual	64	Engineer	1 Adult
P27	Manual	91	Retired	1 Adult
P28	Direct L.	73	Retired	1 Adult
P29	Manual	75	Retired	Na
P30	Indirect L.	28	PhD Student	1 Ad. , 1 Ch.

Internet connection and a central heating control, based on a first-come first-served basis.

Participants were assigned to three groups, each corresponding to one thermostat design, one by one in the order 1, 2, 3. An online budget of £100 was then allocated to each household and participants started to use our system for heating their house. Throughout the study, on each day, their daily heating cost was calculated based on the number of hours their boiler was on (i.e., when the setpoint was higher than the indoor temperature) and the energy prices at those hours. The daily heating cost then was subtracted from their online budget on each day. After four weeks, when the study ended, they received the amount left in their budget in cash as an experimental reward. By so doing, we aimed to encourage participants to respond to the prices, and make savings have a real impact. Similar to the Tariff Agent studies, we used monetary incentives to mimic a real-time pricing scenario.

The simulated real-time pricing scenario was based on actual historical spot prices in the

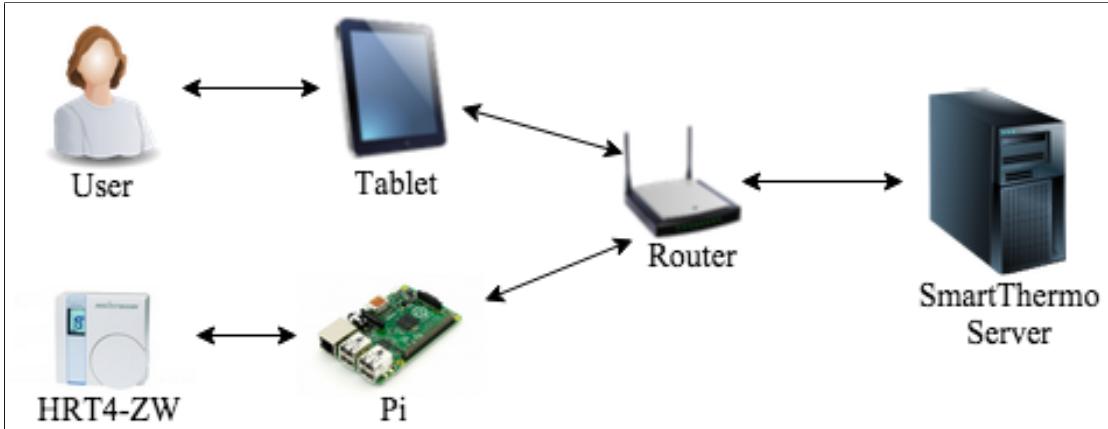


Figure 4.1: Overall system diagram.

UK electricity market in January 2014.² For convenience, we removed extreme outliers from the historical pricing data making the prices range from 5 pence to 35 pence (see Appendix A). During our field study the energy price was changed every 30 minutes, similarly to the UK market.

4.2 Technology

We equipped each household with a Horstmann HRT4-ZW thermostat, a Raspberry Pi (RPi) and an Android 4.4 tablet. Figure 4.1 shows the connections among different entities. The Horstmann thermostat is a standard room thermostat but can be wirelessly controlled over the Z-Wave communication protocol from the RPi (through a RaZberry daughter card³). The RPi also connects through the home wireless broadband router to our web server, where the smart thermostat algorithm and UIs run. The RPi regularly pulls the indoor temperature from the thermostat (every 5 minutes), sends the temperature data to our server, and receives the latest individual heating plan. Based on the plan the RPi then controls the setpoint of the thermostat. The tablet allows participants to access our web application through the broadband connection, and to manipulate their own heating plan. Each tablet was installed with a software called Kiosk Browser Lockdown and our web application was set as the default one. We also added the application as a bookmark on the home screen of participants own devices (tablet or smart phone), if they wished.

²For practicality we recorded the prices about a year earlier than our study took place, that is in February-March 2015.

³<http://razberry.z-wave.me/>

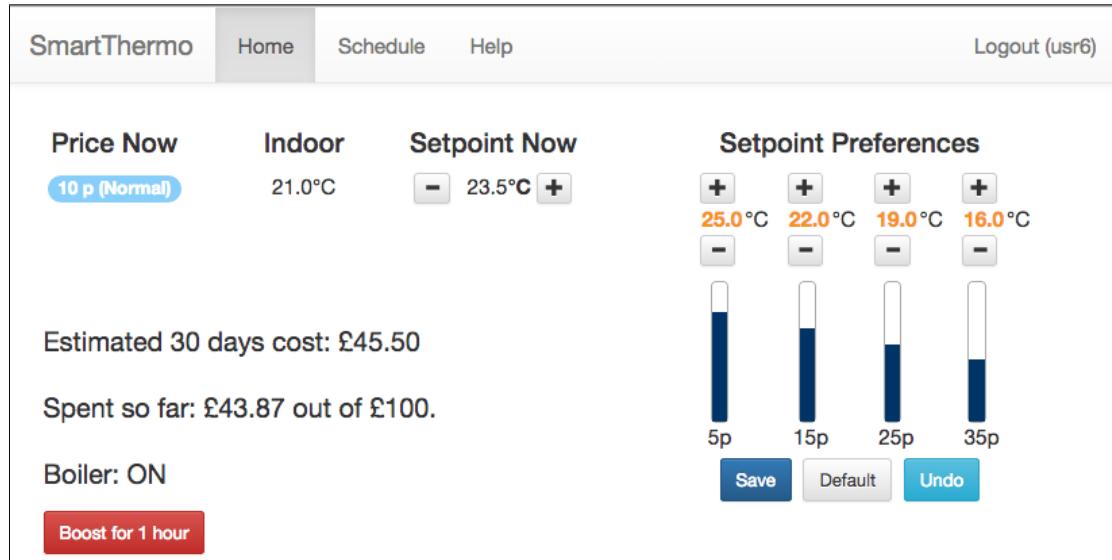


Figure 4.2: Manual - home page.

4.3 Design Variations

In order to gain an understanding of how to best design a smart thermostat for real-time prices, we decided to explore three thermostat designs that we describe in what follows.

4.3.1 Manual Thermostat

This design aims to provide manual operation and involves no machine learning algorithm. Hence, in this design, users are required to manually specify how the temperature is going to be set at different prices through adjusting a number of setpoint sliders. On the home page (Figure 4.2), users can see the current energy price and indoor temperature, and adjust the setpoint by pressing the *+/– buttons* next to it. Each press increases/decreases the setpoint by half a degree. To provide context for the current value, a label indicates whether the price is *normal* (bottom of the range), *high* (mid-range) or *very high* (top of the range). The price value and the label are color coded green, yellow or red for emphasis. At the bottom left of the home page users can find the *boost button*, which allows them to turn the heating on continuously for 1 hour, temporarily overriding the setpoint.

On the right side of the page, four *setpoint sliders* enable users to specify how the setpoint should be changed at different prices.⁴ In other words, these sliders allow users to directly specify how to trade off comfort and cost. These are positioned on the home page to make them easily visible and accessible, even at the risk of increasing the complexity of the page. It should be noted that the sliders are not influenced by the boost button. The sliders are influenced by the *+/–* buttons, and are constrained to

⁴We call these sliders interchangeably setpoint or preference sliders.

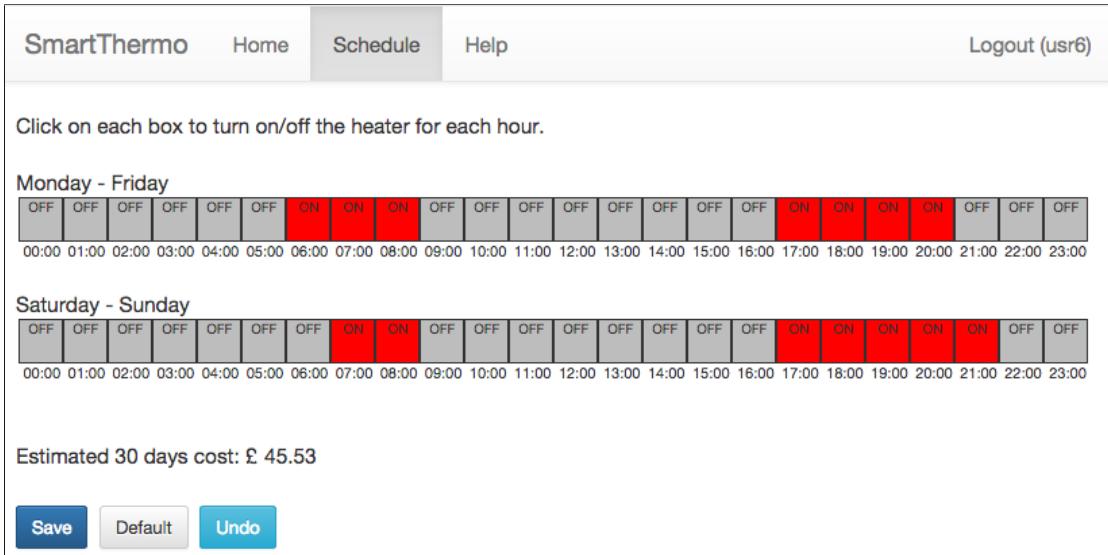


Figure 4.3: Manual - schedule page.

always form a straight line. Thus, if the user changes the setpoint at any slider, the other sliders might change their values as well to conform to the linearity. The temperatures of the intermediate price values are calculated based on linear interpolation.

The schedule page (Figure 4.3) is another page that the manual thermostat users could access. This page allows the users to program the heating schedule that defines the boiler’s on and off times. Due to the screen size of our tablets we decided to divide the schedule of a day into hourly-based time slots and group the days as weekdays and weekend. To change the boiler’s status for a period of time the user only needs to touch on the time slots corresponding to the period. We provided this schedule page since we anticipated that users would expect such a functionality from a smart thermostat.

Both the home and the schedule pages display the ‘Estimated 30 days cost’ that reflects how the current settings on the setpoint sliders and the schedule impact the monthly cost of heating. When users make a change in the sliders or in the schedule, the estimated cost updates accordingly. Also it is important to note that users need to save any changes that they make in order to register that change into the system.

4.3.2 Direct Learning Thermostat

This design uses the machine learning algorithm introduced in a prior work (Shann and Seuken, 2013),⁵ and aims to automate users’ temperature decisions for different prices (see Appendix B for details of how the algorithm works). When users make changes to the temperature, the learning algorithm correlates these changes with the prices and generates a user model. Thus, rather than requiring the user to manually specify the

⁵There might be more advanced algorithms giving better results. We chose this algorithm due to its simplicity and robustness.

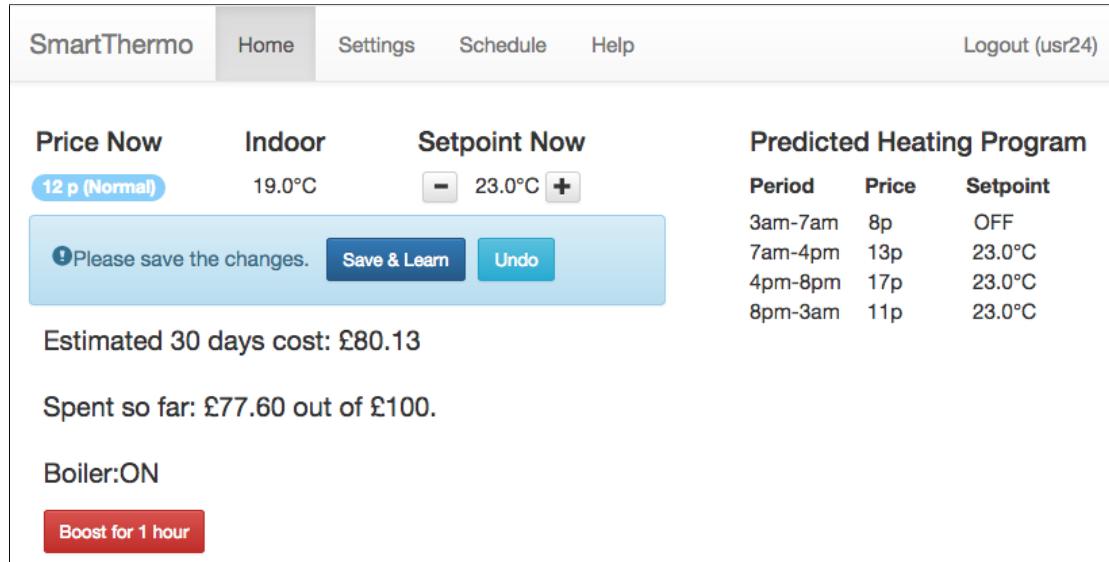


Figure 4.4: Direct learning - home page.

setpoint sliders, the learning algorithm automatically arranges them. Each time the user submits a temperature, the algorithm updates the user's model and the thermostat directly heats to the optimal temperature of the user model based on the current price. The aim of this design is to help users understand that the setpoints that they save are being learned by the smart thermostat for future use to determine the setpoint based on varying prices.

In this thermostat design (Figure 4.4), users directly interact with the machine learning algorithm. When the user presses the $+$ / $-$ buttons, the learning algorithm updates the user's model and displays the optimal setpoint based on the model. The algorithm uses Bayesian inference to update the model, which means it considers the user's individual temperature inputs as noisy data. Thus, the user might need to press the $+$ / $-$ buttons several times to change the setpoint a half degree, depending on the model's prior knowledge. For example, assume that the current setpoint is 18.5°C . If the user now presses the warmer button once, the algorithm will take 19°C as input and do a Bayesian update, resulting in a learned optimal setpoint of 18.7°C , which is then rounded to 18.5°C (the granularity is in steps of 0.5°C). Thus, the user does not see any change in the setpoint. However, if he presses a second time, the algorithm will take 19.5°C as input and the learned optimal setpoint increases to 18.9°C , which will then result in a setpoint change to 19°C . Thus, in this hypothetical example, the user had to press the $+$ button twice to increase the setpoint from 18.5°C to 19°C .

However, to visualise the impact of each press, a pop-up message appears with two buttons, by which users can save or undo the setpoint change. Additionally, each press synchronously affects the 'Estimated 30 days cost' as well as a table called 'Predicted Heating Program'. This table shows the average temperatures that will be set by the thermostat, based on the predicted average prices at four time periods. The setpoint

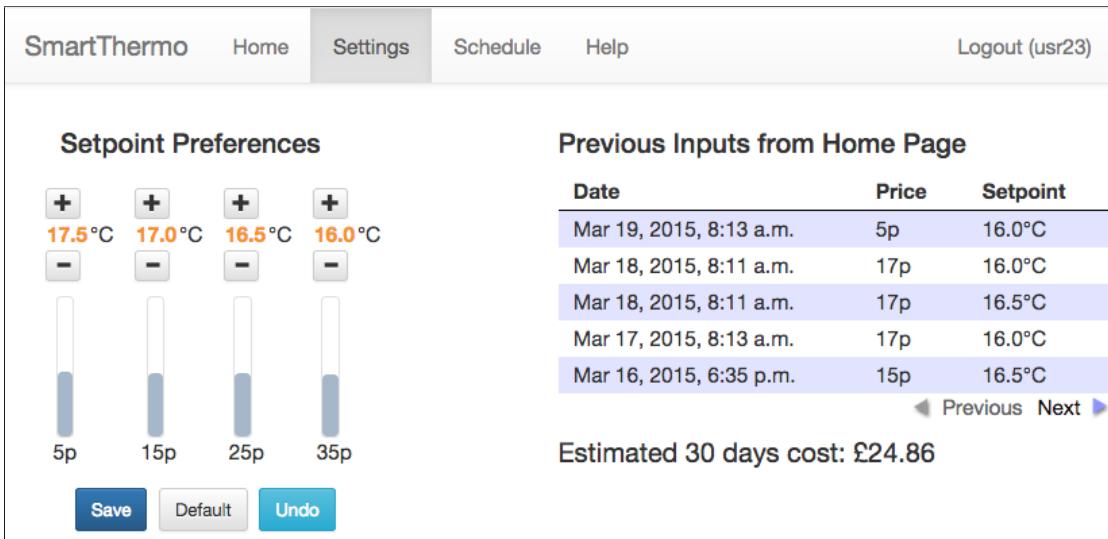


Figure 4.5: Direct learning - settings page.

displayed for each time period changes dynamically according to the updates in the user model. Similar to the manual thermostat, the home page also contains the boost button, which turns the heating on continuously for 1 hour. The boost button does not influence the learning: it was designed as a way to define exceptions to the preferences.

With this thermostat design, users also have an additional page called settings. The settings page (Figure 4.5) aims to provide users an additional level of control and transparency of the learning algorithm. Similar to the home page of the manual design, there are four sliders representing the user's learnt temperature preferences for each price band. We moved the sliders into the settings page because the focus of the learning thermostat is simplicity of use. The user can see how these sliders are arranged by the thermostat by looking at a history table. The table lists the user's previous temperature inputs together with correlated prices. The user can adjust the sliders to specify his own preferences. By so doing, the user resets the learning algorithm and clears the table of previous inputs. Therefore, a confirmation pop-up is shown before the user saves any changes made in the sliders. The schedule page provided in the manual thermostat is also accessible by the users of the direct learning thermostat.

4.3.3 Indirect Learning Thermostat

Similar to the direct learning thermostat, the same learning algorithm is used in this design. However, the rationale behind this design is to enable users to temporarily override the learning, and in this way hide from the users the complexity of the algorithm. Thus, in this design, each time the user submits a temperature, the algorithm updates the user model but the thermostat first heats to the inputted temperature - rather than heating to the optimal temperature of the model. However, after one hour it goes back to auto-mode and sets the setpoint to the optimal temperature of the user model based

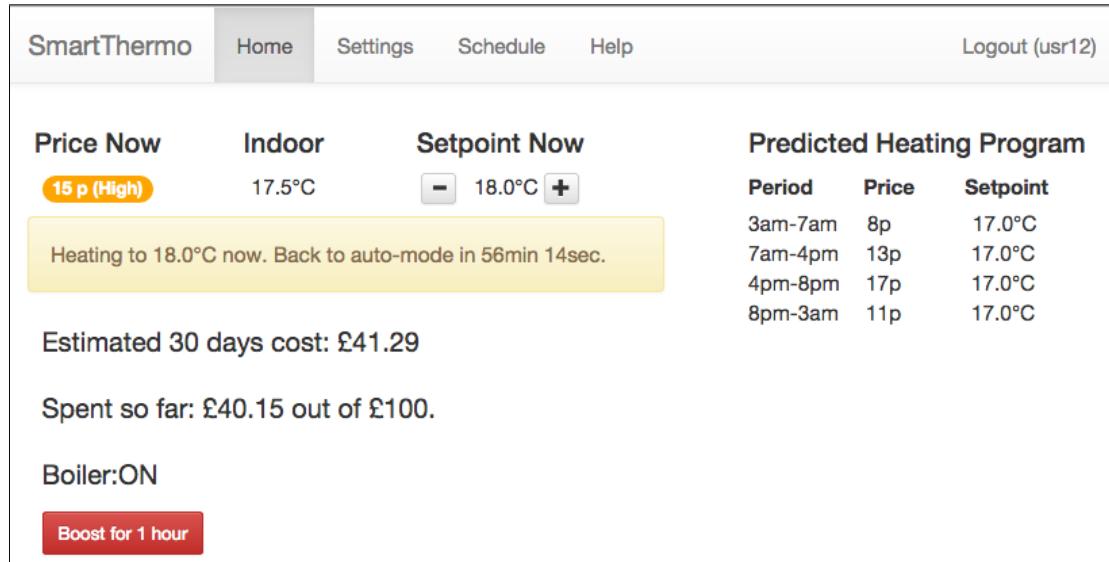


Figure 4.6: Indirect learning - home page.

on the then current price. For example, when the user sets the temperature to 20°C, the thermostat will heat to this exact temperature for one hour. In the background, it takes the 20°C as a new learning input and performs a Bayesian update. After the one hour, the thermostat will switch to the temperature that will be optimal (according to its updated user model) at the then current price.

In this design, since the temperature input that users provide temporarily overrides the setpoint that the algorithm would set based on the user's model and the current price, the +/− buttons work exactly in the same way as in the manual thermostat (each press increases/decreases the setpoint by half a degree). Once the user saves the new setpoint, the algorithm updates the user model based on the new input. Then, it waits for an hour to take the control back and change the setpoint to the learned one according to the then current price. Meanwhile, the thermostat heats to the inputted temperature. This process was explained with a pop-up message including a countdown timer, starting from 60 minutes, in the UI (Figure 4.6). Users can still change the setpoint and save it before the countdown finishes, which will restart the count down with the new setpoint. As in the direct learning thermostat, the home page also contains a boost button, which does not influence the learning. The settings and schedule pages are also provided for the indirect learning thermostats.

4.4 Data Collection and Analysis

During the course of the study, we recorded all users' interactions with the thermostat application. For instance, we recorded when participants changed the setpoint or the heating schedule, or when they used the boost button. Additionally, we collected detailed

quantitative data about the heating habits of each household, including the temperatures the users set in response to real-time prices and how indoor temperatures varied over the course of the study. However, it was difficult to derive conclusions about the impact of the different thermostat designs on these data, since there are other factors affecting people's heating preferences (e.g., weather and home insulation) (Peffer et al., 2011). Therefore, we focused on the qualitative data collected during the interviews.

4.4.1 Interviews

We conducted semi-structured exit interviews with family members at their homes. We interviewed 26 households.⁶ The interviews were mostly held with the participant that signed the consent form at the beginning of the study, however some interviews also involved the participant's partner. In the interviews we asked participants open questions about their use, adoption and understanding of the thermostat. All interviews were audio-recorded, and lasted on average of 34 minutes (SD: 8 minutes, min: 18 minutes, max: 52 minutes).

4.4.2 Analysis

The interviews were fully transcribed and analysed through thematic analysis Braun and Clarke (2006). Four researchers were involved in this, while the coding was performed by two researchers. The analysis started by categorising the material at the sentence level through open codes. Initially 93 open codes were used, later grouped in broader categories that we discuss in the following section.

4.5 Findings

In this section we first present an overview of the quantitative analysis we performed on the overall system usage, based on the automatic interaction logs. Secondly, we report the major findings of our thematic analysis. The analysis revealed six key themes: (1) orientation towards the thermostat's agency, (2) reactions to different UI features, (3) managing the home heating with real-time prices, (4) mental models of the thermostat's learning feature, (5) balancing cost and thermal comfort, and (6) limitations of the thermostat's learning model. We present the categories that revealed these themes in the following sections. In excerpts, we use "F" for female and "M" for male to denote the gender of the household member.

⁶Other 4 participants were not available for the interview.

Table 4.2: Overall quantitative data analysis.

	Manual		Direct L.		Indirect L	
	M	SD	M	SD	M	SD
Setpoint Changes from Home Page	36.7	52.8	13.3	14.4	18.5	22.8
Setpoint Changes from Settings Page	-	-	6.4	7.4	2.5	3
Schedule Changes	18.2	14.8	47.8	51.3	21.8	19.4
Boost Activations	17.5	14.1	7.3	4.3	14.2	11.3
Spent from Budget	£55	£22	£55	£33	£30	£15
Demand Response	36%	19%	34%	17%	47%	24%

Note. M = Mean. SD = Standard Deviation.

4.5.1 Overview of Quantitative Analysis

Users of each thermostat design heated their home using our system for a month. They mostly interacted with the system via the tablet we provided or the tablet they already had. Some of these participants additionally used their mobile phones to access the system. Table 4.2 includes the overall data of system usage. We performed one-way ANOVA tests on the quantitative data. However, we could not find any significant differences or long-term effects in user interactions across the three deployed designs. This was also the case for the analysis of other quantitative data collected (i.e., setpoints and indoor temperatures). The analysis did not reveal any significant differences in users' demand-responses or savings, which might be understandable given the interpersonal, contextual, and environmental differences of the users.

4.5.2 Orientation towards the Thermostat's Agency

All participants commented that they were happy with the thermostat autonomously responding to real-time prices on their behalf. The following is a typical response that we received in the interviews, when we asked participants about their feelings towards the agency of the system.

P18 (M): “I’m happy with that if the thermostat understands that at this price I would rather avoid heating the house, and at this price I would like to heat the house, then I’m happy for it to take over that control, as long as it’s very straightforward for me to override.”

We observed that they felt in control overall and were also mostly confident with the way the thermostat was working.

Analyst: “To what degree did you feel like the system worked for you, or it required you to do the work?”

P13 (F): “So basically I just order the system to do the things for me and the system does the whole thing.”

4.5.3 Reactions to Different UI Features

In the interviews we observed that most participants understood well how to use the UI elements (e.g., the $+$ / $-$ buttons) of each thermostat design. Nearly all participants commented that the thermostat was easy to understand and use. All participants appeared to understand the functionality of the setpoint sliders and mostly appreciated their use.

The users of the direct learning thermostat were mostly aware of the fact that sometimes they were required to press the $+$ / $-$ buttons multiple times to achieve the desired setpoint value. However, they did not explicitly state that this was due to the learning feature.

P3 (F): “The estimated cost would change before the degree thing changed. So, you press it. Sort of, like, Wow it needs four presses per half degree or something, and I was like, because I could see this number here was changed. It was doing something. Then I thought it must be incremental, must be in tenths rather than in halves.”

The indirect learning thermostat users mostly explained the way the 60 minutes countdown works as though it goes back to the previously saved setpoint rather than the learned setpoint based on the price at the time after 60 minutes.

P23 (F): “So then there’s a countdown for 60 minutes and after that the temperature will resume to what was set previously. Occasionally I would reset the temperature again within those 60 minutes.”

4.5.4 Managing the Home Heating with Real-time Prices

Here we detail how and why participants used the thermostat in different ways to heat their house with real-time prices.

4.5.4.1 Setpoint Preferences

Most participants of all three thermostat designs reported preferring to change ‘setpoint now’ from home page to control the indoor temperature with real-time prices. They also fiddled with the setpoint sliders, but the number of times was relatively low compared

to the changes made in the ‘setpoint now’. Interviews revealed that most participants were happy with the arrangement of the sliders and therefore they did not feel the need to alter them often. Also a few participants found the sliders complex, which led them to play with the ‘setpoint now’ more.

We had constrained the setpoint sliders to present a straight line corresponding to users’ heating preferences. Therefore, changes made to one of the sliders affected the values of others. Overall the participants who interacted with the sliders found them easy to adjust and appreciated their use. In the interviews, only one participant griped about this linear relationship among the sliders. However, we noticed that this user only played with the first slider throughout the study, therefore they could make only parallel shift on the slider values without being able to change the slope of the line.

P15 (F): “If we could’ve adjusted them differently and made our own decisions on these rather than they just go up automatically when you change one of the others, we would’ve preferred that.”

Most participants kept the configuration of the sliders in descending order, starting with higher temperature at lower price and lowering the setpoint as the price increases. Two users calibrated the four sliders to have the same setpoint. In other words, they opted for a specific temperature setpoint to heat their house over the course of the study regardless of the heating cost.

P29 (M): “You should be prepared to pay more, a higher rate, if you wanted to be more comfortable. I could change it as when I wanted it, but if I wanted to go to a higher temperature, it could cost me more. But, because I had set it at a flat rate, I wasn’t bothered.”

Also, participants reported that they did not need to change the sliders once they found their limit for how much comfort they could sacrifice to save money. The process of finding such a limit was generally a matter of trial and error:

P11 (F): “It was freezing cold and it must’ve broken. I checked and that’s when I saw it was on 35p and that’s when I changed the lowest set point.”

Some participants were more conscious of and certain about their tolerance limit for temperature even in the early stages of the field study. Therefore, once these participants arranged the sliders early on in the study, they stopped interacting with the sliders, and used other interface features, such as the schedule, to adjust the heating.

P16 (M): “When we first got it, we looked at the pricing bands and made some decisions at that stage. We did it once and I don’t think we revisited it. What we did visit, then, pretty regularly, probably every day, and maybe more often than once a day, we did revisit the schedule.”

4.5.4.2 Schedule

Most participants told us how it was easier to access and change the heating program through our system compared to their previous programmable heating controls. Being able to easily turn on and off the heating by touching on the displayed hourly time slots, and being able to have different programs for weekdays and weekends, seemed to meet participant’s favour.

P21 (F): “We changed it most days because it was so easy to access. If I was going out and I knew that we wouldn’t be home until five, I’d set it to come on at four. Whereas previously, we wouldn’t even touch it on a normal one.”

Here she refers to the “normal one” as a wall-mounted programmable thermostat that she had before taking part in our study. Her comment suggests that people may engage more with their heating systems when the systems are easy to control and access. Specifically, among the users of all UI designs, most participants liked being able to control the heating remotely from anywhere in the house or anywhere outside via their Internet-connected devices - rather than walking to a wall-mounted thermostat each time.

P23 (F): “It was just useful to be able to change the temperature from wherever I was really. I could do it from work and quite often did or if I went from work to the supermarket and then came home, you could do it from the supermarket. So yes, it was really *clever*.”

In the interviews, we observed that occupancy was the major factor affecting the way participants modified the schedule. They mostly tended to turn off the heating when no one was around, and turn it on if someone was at home. Participants who had regular lifestyles reported that they didn’t need to change the schedule often, whereas some participants altered the schedule quite often due to their irregular lifestyles.

P28 (M): “I did change this one [schedule] at the very beginning, but other than that I haven’t touched it at all because, I’ve been working with this time limits for the boiler on and off for 25 years. It’s suited my lifestyle. I’m a creature of habit really.”

P10 (F): "I'm retired so could be at home all day but at the last minute suddenly go off somewhere and we're three adults. So it's three people leading separate lives in a way rather than if we were a family with children and you'd know you would be in the house until half past eight go to school picking up. So our lifestyle is quite *erratic*."

We also observed that some participants used the schedule as a medium to respond to changing energy prices. For example, in the following quote, P9 (F) indicates that she moved the time that the thermostat normally comes on in the mornings nearly one hour earlier to benefit from lower prices.

Analyst: "How do you feel about the real-time prices for heating energy?"

P9 (F): "I noticed that it [price] was cheaper before 7am. Previously, I'd been putting the heating on like maybe 6:45 because we get up about 7:00, then leaving it on while we're getting ready for work and school and then turning it off. I changed that and started putting it on earlier, putting it on at 6:00 and then having it go off at 7:00, and it still kept the house warm enough until we went out sort of an hour or so later."

There were other factors that influenced participants' heating program, such as their daily activities or weather conditions. Participants mostly turned the heating on at times that they usually took showers, or turned it off when they used the oven. While on cold days participants arranged the schedule to make the thermostat come on for more time slots, they had fewer time slots on for milder days.

4.5.4.3 Boost

After deciding how to balance comfort and cost, participants tended to use the "boost" button for exceptional situations to turn the heating on instead of changing the setpoints on the sliders.

P2 (M): "I tended to be comfortable with the settings that I had on it and sort of left it. The only time if it was really cold in the mornings when we got up, I'd press the boost to boost it and it probably only went on for an hour or two."

Some participants also commented that they used the boost button just to heat their home a bit more when the prices were lower.

P21 (F): "I liked when it said £0.05 and I was like; yes, put the heating on, boost it!"

4.5.5 Mental Models of the Thermostat's Learning Feature

Only the users of direct and indirect thermostats were exposed to the machine learning algorithm. These users were required to click the save and learn button that appears every time they make a change in the setpoint from home page in order to register their preferred temperature into the thermostat. Including the text 'learn' in the save button seemed to be successful at conveying the fact that the thermostat was learning. However, when we asked the users' opinion about what the thermostat was learning in the interviews, three users reported that they had not thought about it before and therefore that they had no comment.

Among the participants who formed opinions about the learning feature, most participants appeared to have an understanding that is well-matched with the actual underpinnings of the thermostat's learning feature. Most participants were aware of the fact that the thermostat was trying to correlate their preferred setpoints to varying prices. It seemed that the display of previous inputs in the settings page supported their comprehension. Though, conceivably, no one seemed to be interested in how the thermostat was actually calculating the setpoint based on their previous inputs

Analyst: "If you had to explain the learning feature to one of your friends, how would you explain it?"

P3 (F): "It [thermostat] learns your tolerance for an increase in price. It learns your habits and your behaviours in terms of the price versus the temperature, and then it applies those, reapplies them for future events when the unit price goes up."

P23 (M): "As you input your set point changes according to the prices and then the system starts to understand what your views are of that cost I suppose. That is what you think is expensive and that is what you think is cheap, and then make changes."

In these quotes, the participants are very clear about what the thermostat was trying to learn. They explain that the thermostat was learning their temperature preferences for different prices based on their previous inputs. Also, the participants express that the thermostat was learning in order to be able to autonomously respond to the changing prices on their behalf.

On the other hand, some participants had another interesting mental model description of the learning feature, which was neglecting the effect of the prices. These participants described the thermostat as though it was matching their preferred temperatures with the times and the days of the previous temperature inputs that they provided.

Analyst: "So can you tell me what happens when you click to the save and learn button after you change the setpoint?"

P21 (M): "Well, it [thermostat] updates and it changes the kind of the set-point to what it is going to heat to, but it also learns what you have done. So, I am guessing that later on, if you are doing that at a certain point every day then it's going to learn that."

P30 (M): "If I play a particular temperature as the setpoint and then click on save and learn, from what I understand is the system will take this reading to consideration for whether to turn the boiler on or off but at the same time try to see that at this particular time of the day, whether it's weekday or weekend and then try to replicate that during other days."

This misinterpretation of the learning feature was more prevalent among the indirect learning thermostat users than among the users of the direct learning thermostat. Further exploration also revealed that none of the participants having the misinterpretation was familiar with the Nest thermostat, or in fact any other smart thermostats that exist in the energy market. We can therefore assume that they were not biased.

4.5.6 Balancing Cost and Thermal Comfort

Even though seeing the current price of energy had mostly impacted on how our participants heated their home over the course of the study, there were other significant factors that played key roles in the decisions of the participants for maintaining their thermal comfort at home. One of these key factors was occupancy. Most of our participants commented that they tended to turn off the heating for the times that no one was at home. Another important factor was outside weather as opposed to the indoor temperature: the colder the weather was, the longer the heating was on. Lastly, daily activities at home seemed to substantially influence the participants' heating preferences. While sitting still or having a shower caused participants to turn on the heating, cooking or other physical activities led them to keep the heating off.

P2 (M): "I generally go out by nine o'clock I had the heating going off at eight o'clock in the morning. So it sort of warmed us up to have our showers and be comfortable in the morning, and weekends it depended whether we were in or out as to whether we left it on or knocked it off. So it revolved around our lifestyle and work patterns and things. And the temperature outside. If it was really cold outside then we would have it on longer."

P1 (M): "Most of the time I tried to connect the schedule with my daily activities. For instance, I take shower in the morning, and sometimes I work here at home between 9 and 11. So these are the times that I turn on the heating. Most of the time between 12 and 3, I cook and turn off the heating, because it really doesn't feel cold."

Another interesting finding that emerged from the interviews was the ways participants attempted to maintain their thermal comfort at home without using our heating system. The most prevalent attempt was putting on one more layer of clothing (generally a jumper), or using a blanket when the energy prices are high. Also some participants took the advantage of their other heating sources such as wood-burning stoves, which is typical in small town houses in the UK.

P8 (M): "I think we have probably spent less on our heating in general than we would have done normally. Normally we heat the house pretty much all the time in the winter. We did at times just put another jumper on."

4.5.7 Limitations of the Thermostat's Learning Model

As it is clear in the previous excerpts, the price was not the only factor affecting our participant's setpoint preferences. However, the learning algorithm used in both direct and indirect learning thermostats was only considering two inputs: the setpoint registered and the price at that current time. Therefore, the thermostat was automating the setpoint control only based on the price. This limited learning capability resulted in dissatisfaction among a few participants since the setpoint automatically set by the thermostat was not always the right temperature for its owner. The following quotes are the only ones from which we received such feedback from the participants.

P3 (F): "There were times when I came in and I was like. Hang on a sec. My house is really warm and it must have been because it had learned something that. To do with the temperature. So, it must have said all the prices are this, so they like it warm when it's like this. It's like. Hell no. It's too hot!"

P9 (F): "Well, if I understood the intention that it was trying to set my temperature according to the price, that didn't really work for me. I kind of wanted a combination. I kind of could see the point of that. But like I said, at night, I didn't want it so warm, though perhaps I quite sort of would like it to keep it a degree or two cooler when the temperature's high to save money or something like that. But I also wanted it to let me decide more and not decide for me all the time."

4.6 Discussion

In this section, we revisit the major findings of our study, and discuss them in light of prior literature. We also present implications for interaction design of smart energy systems and for future research.

4.6.1 Designing a Thermostat for Real-Time Prices

Any thermostat designed for real-time prices will need to automate the heating at some level, as otherwise it would be a very difficult task for a human to monitor every price change and alter the heating accordingly. In Chapter 3, we suggested that autonomous systems should allow their users to easily override the automated decisions at any point in time, without completely disabling the system’s autonomy. In this vein, in the Smart Thermo study we observed that some participants used the boost button as a means to temporarily override their temperature preferences for exceptional situations, rather than resetting the learned preferences. These exceptional situations not only occurred when users felt cold and wanted to heat the house despite the high prices, but also happened when users wanted to heat the house a bit more than they would do normally in order to benefit from low prices (typically termed the rebound effect).

One of the most-liked features was the display of ‘Estimated 30 days cost’. As P28-m said, “I’ve watched also my estimated cost each day, to see whether it varied at all. I had taken an interest in it, every day really, I’ve become almost fixated by it.” We observed that the participants used it as a ‘sandbox’ area (Mennicken et al., 2014), by which they could view the consequences of different settings on the cost before approving them. Another well-liked feature was the ability to control the thermostat remotely. Participants commented that this feature affords them a high degree of convenience for heating their home. Most of them reported that they monitored their house (whether the heating was on or off) while they were away, or turned the heating on just before coming home. Furthermore, they found the use of it handy even within the house. For instance, one of our participants commented that she liked being able to take the tablet with her to bed so that she could turn the heating on in cold mornings without having to leave the bed.

Most participants found the thermostat’s heating schedule easy to access and program. However, some participants perceived its hourly time slots as limiting their scheduling plan. This is understandable when one considers that today most heating controls provide finer resolutions (e.g., 10 to 30 minutes). Additionally, grouping the daily heating program by weekdays and weekend was not convenient for all participants to accommodate their occupancy patterns. As an example, one participant said her Saturdays and Sundays are totally different. We also had some participants who did not have any occupancy patterns at all and had to adjust the schedule quite a few times in a day.

Therefore, further research is needed to address how to best design heating programs for people with unpredictable lifestyles.

Our field study showed that participants could use our thermostats to effectively manage their home heating and create temperature preferences based on real-time prices. As we expected these temperature preferences varied for different individuals. While most participants set lower temperatures at peak prices compared to lower price periods, reducing the average energy consumption during peak hours by 38% (see Table 4.2), two households kept the same temperature for all price bands. Furthermore, our participants adopted different strategies to respond to real-time prices. While most participants used the setpoint and the setpoint sliders for reacting to changing prices, some participants interestingly used the boost and the schedule features more than adjusting the setpoint for heating their home with real-time prices. This is in line with a previous study that examined people's use and mental model of their heating system (Revell and Stanton, 2014), and revealed that setpoint adjustment was less prevalent among their participants compared to the adjustments of other devices, such as the programmer, override button and radiator valves.

Finally, we noted the several ways that our participants used to maintain their thermal comfort, especially when the prices were high, without using our heating systems, such as putting on one more layer of clothing or using a blanket. These observations show similarity to the findings of previous work (Clear et al., 2013, 2014), which examine students' daily heating habits and report the similar activities without any financial benefits.

4.6.2 Expectations from Smart Home Heating Systems

While most participants perceived the thermostat as "smart" because of its learning capability of preferred temperatures and its ability to automate home heating based on changing prices, for some participants it was enough to describe the thermostat as smart just because of its remote control capability and its programmable schedule. This perception was due to the fact that these functionalities were mostly new to the participants. More importantly they experienced improvement in their quality of life as these functionalities assisted and facilitated their heating task. This finding is in line with prior research suggesting that computing technologies would be perceived to be "smart" if they offer an advantage for the users' daily tasks (Mennicken et al., 2014).

Regarding the learning feature of the thermostats, participants had different explanations and mental models. While some participants described the system as one that was trying to match their temperature preferences with changing prices, other participants thought that the system was learning the times of days that they set temperatures.

Interestingly, participants who used the direct learning thermostat and had no technical background (e.g., P3, an antiques dealer) described more accurate mental models compared to the participants of the indirect learning thermostat with more technical background (e.g., P1-m and P30-m, both computer science PhD students). A previous study examining non-technical users' understandings of an intelligent system suggests that people's initial mental models and misconceptions stayed relatively constant over their study (Tullio et al., 2007). Therefore, we asked our participants, who had this misconception, if they were aware of any commercial smart thermostats, such as Nest that learns your schedule, in order to see if they had any initial knowledge that would have affected their mental models. However, they all reported that they had not heard of any smart thermostat before. This then may suggest that exposing users directly to the outcomes of learning algorithms may help users to create better mental models. Furthermore, while showing the correlation between previous temperature inputs and prices supported the users' understanding, a more useful method could be a notification system that periodically states what has been learned by the system. We believe these results highlight an important implication for future research in interaction with "smart" energy systems to try and discover the source of people's mental models and learning expectations.

Our system learns users' preferred temperatures at different prices to automate home heating. However, from the interviews, it was clear that the price was not the only factor that our users considered for heating their home. Other key factors were outside weather, occupancy and daily activities within the house. Some participants explicitly stated that the use of the thermostat could be more convenient if it could learn their occupancy patterns. Also, outside weather and the activities that they perform during a day within the house have a significant impact on how people feel the indoor temperature. For instance, most of our participants preferred to have the heating on when they shower and have the heating off when they use their oven or perform physical activities. Therefore, future design of learning thermostats should not only take into account occupancy patterns and outdoor temperatures (Peffer et al., 2011), in addition to people's price preferences, but also people's daily routines (e.g., times that they shower and cook).

4.6.3 Studying Future Smart Energy Systems

In order to let participants experience a future scenario, we prototyped our system based on envisioning (Reeves, 2012). Our scenario depicts an energy market in which consumers can respond to real-time prices by using a smart thermostat that automatically controls heating on their behalf. Participants' statements about their perception and adoption of the smart thermostat indicate that combining experimental reward with a deployed prototype is an effective way to convey a future scenario to participants and

allow them to obtain real life experience, echoing the findings represented in Chapter 3 and a recent study of future scenarios (Costanza et al., 2014). Extending our Tariff Agent studies and the study by Costanza et al. (2014), in Smart Thermo study our participants experienced the autonomous actions of the smart energy system not only through financial incentives, but also through the thermostat's automatic temperature changes. Such changes could directly influence our participants' comfort. Yet, similar to the results of previous studies, our participants mostly felt in control of their heating system and demonstrated a generally positive attitude towards the thermostat. Hence, we believe this finding reinforces those from those previous studies, revealing the potential of future autonomous smart energy systems.

One of the prerequisites for taking part in our study was to have a central heating system with a single boiler. However, we did not define any requirements on the type of thermostat previously installed, such as programmable or non-programmable, digital or analog. Our findings revealed that the type of thermostat familiar to our participants influenced their perception and use of our system. In particular, participants not used to a programmable thermostat focused mostly on the schedule feature of our system, since this was a new and significant feature for them. This circumstance turned out to steer attention away from our primary interest: the ability of our thermostats to automatically react to real-time prices. Hence, future research should take user fragmentation into account in the recruiting process of participants in order to improve the effectiveness of system designs and to obtain more focused results.

4.7 Summary

Smart energy systems that leverage machine learning techniques are increasingly integrated in all aspects of our lives, and they are changing the way that we perform our daily activities. The design of these systems plays a key role in how we adapt to and interact with them. Therefore, we need to better understand how to design user interaction with such systems. To this end, in this chapter, we introduced the design and implementation of three different smart thermostats that automate heating based on users' heating preferences and real-time prices. We presented the evaluation of our designs through a field study, where 30 UK households used our thermostats to heat their homes over a month.

Our findings through thematic analysis show that the participants formed different understandings and expectations of our smart thermostat, and used it in various ways to effectively respond to real-time prices while maintaining their thermal comfort. Based on the findings, we provided a number of design and research implications, specifically for designing future smart thermostats that will assist us in controlling home heating with real-time pricing, and for future intelligent autonomous systems. Specifically, we

observed that exposing users directly to the consequences of machine learning resulted in better user mental models. Also we learnt that taking user fragmentation into account in the recruiting process of participants is important to improve the effectiveness of system designs and to obtain more focused results.

We suggest that future learning thermostats should provide users a means to temporarily override their preferences for exceptional situations and a means to view the financial consequences of different settings before confirming them. Moreover, future design of learning thermostats should not only take into account occupancy patterns, outdoor temperatures, people's price preferences, but also people's daily routines (e.g., times that they shower and cook). These recommendations will assist designers in addressing the challenges highlighted in Chapter 1, and therefore will help them improving user experience with smart energy systems. This in return will enable us to more smoothly integrate these systems into our everyday lives and actually benefit from them.

Chapter 5

Conclusions and Future Work

In this chapter, we summarise the contributions of this thesis and give directions for future work.

5.1 Conclusions

In Chapter 1, we first stated that while there exist many studies that propose algorithms for autonomous agents to control home energy use where, for example, agents automate microstorage or appliance use, there are very few studies that examine the use of agents with real-world deployments. Hence, we then highlighted that there is a significant gap in our understanding of how we should design interactions with agent-based energy systems, especially for the ones that might possibly intrude upon our daily activities. Given this, we represented the aim of this work that is to provide novel design guidelines that improve user interactions with autonomous agent-based home energy management systems, and ascertain to what extent actually users embrace autonomous agents, whether they opt for fully autonomous or semi-autonomous agents as opposed to controlling their preferences without any automation.

To this end, Chapter 2 provided a necessary background of previous research approaches that examine human-agent interaction in miscellaneous genres with various perspectives. We also gave an overview of the forthcoming developments in smart energy systems and existing energy related studies from various disciplines work including agent-based computing, HCI, and social sciences. We then concluded that these approaches have neglected the questions of how software agents should be involved in smart energy infrastructures and how users and these agents should interact with each other, which is what we aim to address in this thesis by representing a number of field evaluations of different designs of agent-based energy systems.

Chapter 3 represents two field studies that exposed participants to a prototyped future energy scenario. Our scenario simulates a situation where households can switch electricity tariff on a daily basis, to try and best match their consumption level. This scenario enabled us to study users' interactions with Tariff Agent, an interactive autonomous system designed to help in managing energy costs, which offers flexible autonomy and detailed information about its operation. The studies were made possible by combining off-the-shelf Internet-connected sensors with Web technology, and with monetary rewards. Our field studies enabled participants to experience an autonomous energy service agent in their everyday lives, form opinions about it and develop strategies to integrate its autonomous operation in their everyday practices. Based on the results of our field studies, we demonstrate that users are, at least in part, ready to use systems like Tariff Agent to manage their energy tariffs. However, the results stress that system designs need to strike a nuanced balance between providing the user with means to monitor system performance and take control when they consider it necessary. We then provided novel design guidelines for implementing mixed-initiative interactions with autonomous energy systems. The design guidelines suggest that designers should provide an easy way for users to receive updates about the status and operation of the agent, enable users to instruct the agent by offering them opportunities to declare their plans and integrate these plans into the agent's operations, and leave the system open to transfer of control by allowing users to adjust the system's level of autonomy. We believe these guidelines, which are derived from real-world deployments, can be also used for developing applications in other domains that involve human-agent interaction.

Chapter 4 represents a field evaluation of a smart thermostat called Smart Thermo, which automates heating based on users' heating temperature preferences over real-time prices. 30 UK households used our thermostats to heat their homes over a month. Our findings through thematic analysis indicate that the participants formed different understandings and expectations of our smart thermostat, and used it in various ways to effectively respond to real-time prices while maintaining their thermal comfort. Furthermore, we show that exposing users directly to the consequences of machine learning may result in better user mental models. Based on the findings, we finally provide a number of design implications, specifically for designing future smart thermostats that will assist us in controlling home heating with real-time pricing, and for other future intelligent autonomous systems of which actions may cause financial consequences for us. In particular, the implications suggest that designers should provide users a means to temporarily override their preferences on a learning-based system for exceptional situations, rather than resetting the learned preferences, and a means to view the financial consequences of different settings before approving them.

Tariff Agent and Smart Thermo were developed by taking into account the prototype requirements stated in Section 1.2.1. In both, the system's autonomous actions were

made configurable, visible, and controllable to users. Moreover, both systems were evaluated with field studies based on the evaluation requirements presented in Section 1.2.2. We linked users' actual electricity consumption with monetary incentives to provide ecological validity and enable them to experience varying energy prices. We attached importance to recruit a fair number of participants having different backgrounds (see Tables 3.1 and 4.1). To better understand user perception of and interaction with our systems, we used both quantitative and qualitative research methods. Based on the quantitative and qualitative analyses, we provided novel design guidelines for developing autonomous domestic energy systems for everyday life use, which helped us to meet our research objectives indicated in Section 1.3. Furthermore, our results enabled us to address the research challenges mentioned in Section 1.4. In particular, we showed that flexible autonomy helps users feeling in control and therefore sustaining their trust in and interaction with the agent. We also showed the intelligibility of agents seems to be improved by providing means for users to direct the agent, and by exposing the users to the consequences of their inputs without any delay.

When taken together, the contributions of the field evaluations of the novel agent-based energy applications that are represented in this thesis advances the state of the art interaction design in the domain of human-agent interaction. We believe our contributions highlight opportunities and show promising directions to design autonomous energy systems in ways that make them useful and acceptable to users in their everyday lives. Our design recommendations will assist designers in improving user experience with autonomous technologies, which in return will help us to more smoothly integrate these technologies into our everyday lives and actually benefit from them.

The field studies represented in this thesis are significant steps towards achieving the goal of designing interactive autonomous systems that meet user satisfaction in the domain of residential energy management. However, we recognise that there are some limitations in our studies, which pave the way for our future work. In the following section, we discuss the limitations and our future research focuses.

5.2 Future Work

The interviews with the users of Tariff Agent and Smart Thermo revealed that some participants perceived these systems more smarter than they actually were. For instance, some Tariff Agent users thought the system was learning their consumption habits over the past days, whilst the system was just using the yesterday's consumption. Similarly, Smart Thermo was only learning people's temperature preferences over varying prices. However, some users believed that the thermostat was also learning their temperature preferences based on the times of the day. We believe these results highlight an important

implication for future research in interaction with autonomous agent-based systems to explore the underlying source of people's understandings and learning expectations.

As with all field experiments, our studies are subject to several limitations. These limitations suggest a number of directions for future research. In our studies only one participant per household took part in the interview. Therefore, the focus was on individual interaction with and perception of the system. However, as most of our participants live with others, the system may affect the social dynamics in the home around energy consumption. Further research is required to better understand the potential social impacts of the system on home dwellers, and consequently draw implications on how domestic autonomous systems may best be designed to support multi-user interaction. Moreover, the duration of our field studies were limited to 2, 4 and 6 weeks. An opportunity for future work, then, is to observe user engagement and potential behavioral change through longer term studies. Furthermore, our participants are not representative of the overall society: as almost all of them are educated to above average levels; so it is important to extend this work to a more general population.

Moreover, our contributions through a number of field studies involved an individual person interacting with a single agent. Future work could aim to advance our research through investigating human-agent interactions in multi-agent systems where an individual person could interact with multiple agents or with other people, directly or through their agent. In order to investigate human-agent interactions in multi-agent systems, an agent-based platform could be deployed for energy exchange among multiple consumers. The platform could initiate the energy exchange protocol introduced in a recent study ([Alam et al., 2015](#)) as a future scenario, where off-grid homes equipped with renewable energy sources and electricity storages can negotiate and exchange energy with each other. A field study of such a scenario could help us to better understand how we should design user interaction with multi-agent systems to support people's collaborative activities and their coordination.

Appendix A

Real-Time Energy Prices

The following tables include the real-time prices used in the Smart Thermo study discussed in Chapter 4. The pairs presented in the tables correspond to the actual and predicted prices for each half an hour period of each day, over 31 days, starting from 1st January 2014.

Period/Day	1	2	3	4	5	6	7	8
1	(12,12)	(7,7)	(10,7)	(12,12)	(12,15)	(7,12)	(7,10)	(12,7)
2	(12,12)	(7,7)	(12,12)	(12,12)	(12,15)	(12,12)	(10,10)	(12,7)
3	(12,12)	(7,7)	(12,12)	(10,7)	(10,15)	(12,12)	(12,10)	(12,7)
4	(12,10)	(7,7)	(7,7)	(10,7)	(10,12)	(12,12)	(7,7)	(12,7)
5	(7,10)	(5,7)	(7,7)	(7,7)	(10,12)	(7,12)	(5,7)	(7,7)
6	(7,10)	(5,7)	(7,7)	(5,7)	(10,10)	(7,12)	(5,7)	(12,7)
7	(10,12)	(5,7)	(7,7)	(5,7)	(7,10)	(5,10)	(5,7)	(7,7)
8	(5,12)	(5,7)	(5,7)	(5,7)	(7,10)	(5,7)	(5,7)	(7,7)
9	(5,12)	(5,7)	(5,7)	(5,7)	(7,10)	(5,7)	(5,7)	(7,7)
10	(5,7)	(5,7)	(5,7)	(5,7)	(7,7)	(5,7)	(5,7)	(7,7)
11	(5,7)	(5,7)	(5,7)	(5,7)	(12,7)	(5,7)	(5,7)	(7,7)
12	(5,7)	(5,7)	(5,7)	(10,7)	(10,7)	(5,7)	(12,7)	(7,7)
13	(5,7)	(7,10)	(5,10)	(12,10)	(12,7)	(7,15)	(12,10)	(10,10)
14	(5,7)	(7,10)	(7,10)	(7,10)	(7,7)	(7,7)	(15,10)	(10,15)
15	(5,7)	(7,10)	(7,10)	(7,12)	(7,10)	(10,10)	(15,10)	(12,12)
16	(5,7)	(7,10)	(10,10)	(7,12)	(10,10)	(15,10)	(15,10)	(15,17)
17	(5,5)	(10,12)	(12,10)	(7,10)	(7,10)	(12,10)	(10,10)	(15,12)
18	(5,7)	(15,12)	(12,12)	(10,10)	(7,10)	(10,10)	(10,10)	(10,12)
19	(7,7)	(15,10)	(12,12)	(10,17)	(10,12)	(12,12)	(10,12)	(15,12)
20	(7,7)	(15,12)	(12,12)	(10,22)	(10,12)	(15,12)	(10,12)	(10,12)
21	(7,7)	(12,12)	(12,10)	(10,20)	(12,12)	(12,12)	(15,12)	(10,12)
22	(7,7)	(12,12)	(12,10)	(10,17)	(10,12)	(12,12)	(10,12)	(10,12)
23	(10,10)	(12,12)	(12,10)	(25,15)	(12,12)	(10,12)	(10,10)	(10,12)
24	(12,10)	(12,12)	(12,12)	(20,12)	(12,12)	(10,12)	(10,10)	(15,12)
25	(15,10)	(12,12)	(12,12)	(27,12)	(15,12)	(10,12)	(10,10)	(15,12)
26	(12,10)	(12,12)	(12,12)	(25,12)	(15,10)	(10,12)	(10,10)	(15,12)
27	(12,10)	(12,12)	(12,12)	(20,12)	(12,10)	(10,12)	(12,10)	(15,15)
28	(12,10)	(12,12)	(15,12)	(17,12)	(10,10)	(10,12)	(12,10)	(17,15)
29	(12,10)	(12,12)	(12,12)	(15,12)	(12,10)	(10,12)	(12,10)	(20,12)
30	(12,10)	(12,10)	(12,10)	(15,15)	(10,10)	(10,12)	(12,10)	(20,12)
31	(10,10)	(10,12)	(12,12)	(15,12)	(10,12)	(10,12)	(10,15)	(25,12)
32	(10,10)	(10,12)	(12,12)	(15,12)	(12,12)	(12,12)	(10,12)	(25,12)
33	(15,12)	(15,15)	(15,15)	(25,15)	(12,12)	(12,20)	(12,15)	(30,12)
34	(20,25)	(25,17)	(20,15)	(35,17)	(15,15)	(20,22)	(22,20)	(35,27)
35	(20,27)	(35,17)	(20,17)	(35,20)	(20,17)	(17,22)	(25,17)	(15,30)
36	(17,25)	(35,17)	(17,15)	(35,20)	(12,17)	(20,22)	(22,17)	(15,27)
37	(17,22)	(30,15)	(15,15)	(35,17)	(20,15)	(17,20)	(20,15)	(15,15)
38	(15,15)	(15,15)	(12,15)	(32,15)	(15,15)	(15,17)	(15,15)	(15,15)
39	(15,15)	(15,12)	(15,12)	(17,12)	(15,12)	(12,15)	(15,12)	(12,12)
40	(15,15)	(15,12)	(15,12)	(15,12)	(10,12)	(12,15)	(15,12)	(15,12)
41	(15,15)	(12,12)	(15,12)	(12,12)	(12,12)	(12,15)	(12,12)	(12,12)
42	(10,12)	(12,12)	(15,12)	(12,12)	(12,12)	(12,12)	(12,12)	(12,12)
43	(10,17)	(10,10)	(15,12)	(15,12)	(12,12)	(10,12)	(10,12)	(12,12)
44	(10,12)	(10,10)	(12,12)	(12,12)	(12,12)	(10,12)	(10,12)	(12,12)
45	(7,12)	(10,10)	(10,10)	(12,12)	(12,12)	(10,12)	(10,10)	(10,12)
46	(7,10)	(10,10)	(10,10)	(12,12)	(12,12)	(7,10)	(10,10)	(10,10)
47	(7,7)	(7,7)	(12,12)	(12,15)	(7,15)	(7,10)	(7,10)	(7,10)
48	(7,7)	(10,7)	(12,15)	(12,15)	(7,12)	(7,7)	(7,10)	(7,10)

Period/Day	9	10	11	12	13	14	15	16
1	(12,10)	(12,15)	(10,12)	(12,10)	(12,12)	(10,10)	(12,10)	(12,10)
2	(12,10)	(15,15)	(12,12)	(12,10)	(12,12)	(12,15)	(12,10)	(12,10)
3	(12,10)	(12,15)	(12,12)	(12,10)	(12,10)	(12,15)	(12,10)	(12,10)
4	(12,10)	(10,12)	(10,10)	(12,10)	(12,10)	(10,15)	(10,10)	(10,12)
5	(7,7)	(10,10)	(10,10)	(12,7)	(12,10)	(10,15)	(10,12)	(10,10)
6	(7,7)	(10,12)	(12,10)	(12,7)	(12,10)	(12,15)	(10,10)	(10,10)
7	(7,12)	(7,12)	(12,7)	(12,7)	(12,15)	(12,10)	(10,10)	(7,10)
8	(5,7)	(7,10)	(7,7)	(12,7)	(7,7)	(10,10)	(10,10)	(7,10)
9	(7,12)	(7,10)	(7,7)	(12,7)	(7,15)	(10,10)	(10,10)	(7,10)
10	(5,7)	(7,10)	(7,7)	(7,7)	(7,10)	(10,10)	(5,7)	(12,10)
11	(7,10)	(10,10)	(7,10)	(7,7)	(7,7)	(10,12)	(10,12)	(7,12)
12	(7,15)	(12,10)	(7,10)	(7,7)	(7,7)	(10,12)	(7,15)	(7,12)
13	(10,15)	(15,12)	(7,12)	(7,7)	(15,7)	(12,12)	(7,15)	(10,12)
14	(10,15)	(15,12)	(7,12)	(7,7)	(15,7)	(12,12)	(10,15)	(15,15)
15	(10,15)	(15,12)	(10,17)	(12,7)	(15,7)	(12,15)	(12,20)	(15,17)
16	(17,20)	(22,12)	(10,25)	(10,10)	(27,7)	(15,17)	(15,30)	(15,17)
17	(30,17)	(15,12)	(7,30)	(10,10)	(17,10)	(15,15)	(12,27)	(27,12)
18	(12,17)	(15,12)	(10,12)	(10,10)	(12,12)	(15,17)	(15,20)	(12,25)
19	(12,15)	(15,12)	(10,12)	(10,12)	(17,15)	(12,25)	(15,17)	(12,30)
20	(12,17)	(15,17)	(10,12)	(12,12)	(20,25)	(12,32)	(15,17)	(12,32)
21	(12,15)	(15,17)	(12,12)	(12,12)	(12,32)	(12,25)	(15,17)	(12,25)
22	(12,15)	(15,15)	(15,17)	(12,12)	(15,17)	(12,22)	(15,17)	(12,27)
23	(12,12)	(15,20)	(12,17)	(12,12)	(12,17)	(10,17)	(15,17)	(12,17)
24	(12,12)	(15,17)	(12,15)	(17,12)	(12,17)	(10,17)	(12,15)	(12,20)
25	(12,15)	(12,17)	(10,15)	(27,12)	(12,17)	(10,17)	(12,15)	(12,20)
26	(10,12)	(12,17)	(10,12)	(30,12)	(12,17)	(10,15)	(12,15)	(12,15)
27	(10,12)	(15,17)	(10,12)	(30,12)	(12,17)	(10,17)	(12,12)	(10,15)
28	(10,12)	(15,15)	(10,12)	(32,12)	(12,17)	(10,15)	(12,12)	(10,15)
29	(10,12)	(17,17)	(10,12)	(32,10)	(12,17)	(10,17)	(12,12)	(12,15)
30	(10,12)	(15,15)	(10,12)	(27,10)	(12,15)	(10,15)	(12,12)	(12,15)
31	(10,12)	(12,15)	(12,12)	(22,10)	(12,12)	(10,12)	(10,12)	(15,15)
32	(10,12)	(15,15)	(12,12)	(17,12)	(15,15)	(10,15)	(12,12)	(20,15)
33	(12,12)	(15,12)	(12,12)	(25,12)	(12,12)	(10,17)	(12,12)	(22,17)
34	(12,25)	(20,17)	(15,15)	(25,12)	(17,25)	(15,20)	(17,17)	(27,32)
35	(15,32)	(17,20)	(25,17)	(27,35)	(20,32)	(15,20)	(20,27)	(27,22)
36	(15,25)	(15,17)	(32,15)	(22,35)	(20,27)	(15,20)	(20,25)	(20,20)
37	(12,27)	(15,17)	(27,12)	(20,30)	(17,15)	(12,20)	(12,20)	(17,20)
38	(12,20)	(15,15)	(25,12)	(17,15)	(15,15)	(12,17)	(15,17)	(15,17)
39	(15,15)	(12,12)	(20,12)	(15,12)	(12,12)	(12,12)	(17,12)	(12,15)
40	(15,15)	(15,12)	(15,12)	(15,12)	(15,15)	(12,12)	(17,12)	(12,15)
41	(15,12)	(12,12)	(12,12)	(15,12)	(12,17)	(15,15)	(10,12)	(12,12)
42	(12,12)	(12,12)	(10,12)	(15,12)	(10,12)	(15,12)	(17,12)	(12,12)
43	(12,15)	(12,12)	(12,12)	(15,10)	(12,15)	(12,15)	(15,12)	(17,12)
44	(10,12)	(12,15)	(12,15)	(10,10)	(12,15)	(10,15)	(10,12)	(12,15)
45	(10,12)	(12,15)	(10,12)	(10,10)	(10,10)	(12,15)	(10,15)	(12,15)
46	(10,12)	(12,12)	(12,12)	(7,10)	(12,10)	(12,12)	(10,12)	(10,15)
47	(12,12)	(12,10)	(12,10)	(7,12)	(12,10)	(10,10)	(10,10)	(10,12)
48	(12,12)	(12,10)	(12,10)	(12,12)	(10,10)	(10,10)	(10,10)	(10,10)

Period/Day	17	18	19	20	21	22	23	24
1	(10,15)	(15,10)	(12,12)	(10,10)	(12,12)	(7,15)	(10,12)	(7,17)
2	(12,15)	(15,10)	(12,12)	(10,15)	(12,12)	(7,15)	(10,15)	(12,17)
3	(12,17)	(15,10)	(12,15)	(10,15)	(10,15)	(7,15)	(12,17)	(12,17)
4	(7,15)	(12,7)	(10,15)	(10,15)	(10,12)	(7,17)	(7,17)	(12,15)
5	(7,15)	(12,7)	(7,15)	(12,15)	(10,12)	(7,15)	(10,15)	(12,15)
6	(12,15)	(12,7)	(7,15)	(12,15)	(10,12)	(7,15)	(10,15)	(12,17)
7	(7,12)	(10,7)	(7,15)	(10,15)	(7,12)	(7,15)	(7,15)	(12,15)
8	(7,12)	(10,7)	(7,7)	(7,10)	(7,12)	(7,15)	(7,15)	(7,15)
9	(7,15)	(7,7)	(7,15)	(7,7)	(7,12)	(7,15)	(7,17)	(7,17)
10	(7,12)	(7,7)	(7,15)	(7,7)	(7,12)	(7,15)	(7,12)	(7,10)
11	(7,12)	(7,7)	(7,15)	(10,7)	(10,12)	(7,12)	(7,15)	(7,15)
12	(7,15)	(7,7)	(7,15)	(10,7)	(7,10)	(7,12)	(7,15)	(7,12)
13	(7,15)	(7,10)	(7,7)	(12,7)	(12,12)	(7,15)	(10,12)	(12,15)
14	(15,15)	(7,12)	(7,7)	(12,7)	(12,12)	(10,15)	(10,15)	(12,15)
15	(15,17)	(7,12)	(7,12)	(12,7)	(12,12)	(10,17)	(10,17)	(12,15)
16	(12,20)	(7,12)	(7,10)	(15,7)	(12,12)	(15,17)	(15,17)	(12,15)
17	(12,17)	(10,12)	(12,10)	(15,10)	(12,12)	(12,12)	(12,17)	(12,12)
18	(12,12)	(10,12)	(12,12)	(12,12)	(12,12)	(12,12)	(12,17)	(12,12)
19	(12,15)	(10,12)	(12,15)	(12,10)	(12,12)	(12,17)	(12,17)	(12,12)
20	(12,17)	(15,17)	(12,17)	(15,15)	(12,15)	(12,22)	(12,17)	(12,12)
21	(12,15)	(10,12)	(15,17)	(15,15)	(12,32)	(12,17)	(12,20)	(15,12)
22	(12,17)	(10,17)	(15,25)	(15,17)	(12,32)	(12,17)	(12,20)	(15,12)
23	(12,15)	(12,17)	(12,22)	(27,17)	(15,27)	(12,15)	(10,17)	(15,12)
24	(12,17)	(12,17)	(12,25)	(27,17)	(15,17)	(10,12)	(10,17)	(15,12)
25	(12,12)	(10,17)	(12,25)	(12,20)	(10,22)	(10,12)	(10,20)	(12,12)
26	(10,12)	(10,15)	(12,20)	(12,17)	(10,17)	(10,12)	(10,20)	(12,12)
27	(10,12)	(10,17)	(12,17)	(12,17)	(15,12)	(10,15)	(12,17)	(12,12)
28	(10,15)	(10,17)	(12,17)	(12,15)	(15,12)	(10,15)	(12,17)	(12,12)
29	(12,17)	(10,10)	(12,17)	(12,15)	(12,15)	(10,15)	(10,15)	(12,15)
30	(10,15)	(7,10)	(12,15)	(12,15)	(10,12)	(10,15)	(10,15)	(12,22)
31	(10,15)	(10,12)	(12,17)	(12,15)	(12,17)	(12,15)	(10,15)	(10,15)
32	(12,15)	(10,12)	(12,17)	(15,15)	(12,17)	(12,12)	(10,17)	(10,15)
33	(10,15)	(10,12)	(12,17)	(17,15)	(15,12)	(12,12)	(10,15)	(10,15)
34	(30,15)	(10,15)	(15,25)	(17,15)	(15,15)	(12,15)	(17,20)	(12,15)
35	(32,20)	(15,15)	(22,32)	(17,22)	(15,17)	(15,15)	(35,35)	(15,20)
36	(27,20)	(15,15)	(22,30)	(17,20)	(15,15)	(12,15)	(32,35)	(12,20)
37	(27,17)	(12,15)	(20,17)	(17,12)	(15,15)	(15,15)	(25,35)	(12,17)
38	(12,17)	(12,15)	(17,15)	(17,15)	(12,15)	(12,15)	(20,27)	(15,17)
39	(17,12)	(10,15)	(15,12)	(15,15)	(15,15)	(12,15)	(15,15)	(15,12)
40	(17,12)	(10,12)	(12,12)	(15,12)	(15,15)	(12,17)	(12,15)	(15,12)
41	(15,12)	(10,12)	(12,12)	(12,12)	(15,15)	(12,15)	(12,15)	(15,12)
42	(15,12)	(12,12)	(12,12)	(12,12)	(12,12)	(12,15)	(12,15)	(15,12)
43	(15,12)	(12,12)	(12,12)	(12,15)	(10,12)	(12,15)	(10,15)	(15,12)
44	(12,12)	(10,12)	(10,12)	(10,12)	(10,12)	(10,15)	(10,15)	(10,12)
45	(12,12)	(7,12)	(10,12)	(10,15)	(10,12)	(10,12)	(10,15)	(10,12)
46	(15,12)	(12,10)	(10,10)	(10,15)	(10,10)	(10,12)	(10,12)	(10,12)
47	(12,10)	(12,10)	(10,12)	(12,12)	(7,10)	(10,10)	(7,15)	(10,10)
48	(12,10)	(12,10)	(10,10)	(12,10)	(7,15)	(10,10)	(7,15)	(10,10)

Period/Day	25	26	27	28	29	30	31
1	(10,10)	(12,17)	(7,10)	(15,7)	(10,7)	(12,12)	(12,15)
2	(10,10)	(12,20)	(7,12)	(12,7)	(10,15)	(12,12)	(12,15)
3	(10,10)	(12,17)	(7,12)	(12,7)	(12,12)	(10,17)	(12,15)
4	(10,10)	(12,15)	(7,12)	(12,7)	(10,12)	(10,15)	(10,15)
5	(7,10)	(12,15)	(7,10)	(12,7)	(10,12)	(10,20)	(10,17)
6	(7,10)	(12,15)	(7,12)	(12,7)	(10,12)	(10,17)	(10,17)
7	(7,10)	(12,15)	(7,7)	(7,7)	(7,12)	(10,12)	(10,17)
8	(7,10)	(10,15)	(7,7)	(7,7)	(7,12)	(10,12)	(10,17)
9	(7,10)	(7,15)	(7,7)	(7,7)	(7,12)	(7,7)	(7,15)
10	(7,10)	(7,15)	(7,7)	(7,7)	(7,12)	(7,7)	(7,15)
11	(7,10)	(7,10)	(7,7)	(7,7)	(7,17)	(15,15)	(10,15)
12	(7,10)	(7,10)	(7,7)	(7,7)	(7,12)	(15,15)	(10,10)
13	(7,10)	(7,15)	(10,7)	(12,15)	(10,17)	(12,25)	(10,15)
14	(7,15)	(7,15)	(12,7)	(12,15)	(12,17)	(12,17)	(10,15)
15	(7,12)	(7,15)	(15,7)	(12,17)	(10,22)	(15,17)	(12,17)
16	(7,12)	(7,17)	(15,7)	(15,17)	(10,27)	(15,17)	(17,17)
17	(15,12)	(7,17)	(15,7)	(12,17)	(12,22)	(17,15)	(17,20)
18	(15,12)	(7,17)	(15,7)	(12,17)	(10,20)	(17,15)	(12,12)
19	(15,12)	(12,17)	(15,10)	(15,12)	(10,20)	(27,15)	(12,17)
20	(15,12)	(15,17)	(15,10)	(15,17)	(10,20)	(27,15)	(12,12)
21	(15,17)	(12,17)	(15,10)	(15,12)	(10,22)	(27,15)	(12,12)
22	(15,17)	(15,17)	(15,10)	(15,12)	(10,17)	(27,15)	(12,12)
23	(15,17)	(15,15)	(15,12)	(15,12)	(10,15)	(27,12)	(15,12)
24	(15,17)	(15,17)	(15,12)	(15,12)	(15,15)	(35,12)	(15,12)
25	(15,15)	(15,17)	(15,12)	(15,20)	(15,15)	(35,12)	(15,12)
26	(15,15)	(15,17)	(15,12)	(15,20)	(15,15)	(30,12)	(15,12)
27	(10,15)	(15,15)	(12,12)	(17,20)	(15,12)	(22,12)	(15,10)
28	(12,15)	(15,15)	(12,10)	(12,17)	(12,12)	(15,12)	(15,10)
29	(12,15)	(15,12)	(15,10)	(12,20)	(12,10)	(17,12)	(15,10)
30	(10,15)	(15,12)	(15,10)	(12,17)	(12,10)	(15,12)	(15,10)
31	(12,12)	(12,10)	(12,12)	(12,15)	(15,12)	(15,12)	(15,12)
32	(12,15)	(10,10)	(15,12)	(12,15)	(15,12)	(15,12)	(15,12)
33	(12,15)	(10,10)	(12,12)	(12,15)	(15,12)	(15,12)	(15,12)
34	(15,22)	(12,12)	(15,15)	(12,27)	(17,27)	(25,15)	(30,15)
35	(22,25)	(25,22)	(25,25)	(20,27)	(15,27)	(22,30)	(27,17)
36	(22,15)	(25,25)	(22,25)	(17,17)	(15,25)	(17,27)	(27,17)
37	(20,15)	(20,15)	(20,22)	(15,20)	(15,22)	(15,25)	(25,15)
38	(15,15)	(17,12)	(20,22)	(12,15)	(15,15)	(15,22)	(22,15)
39	(12,15)	(15,12)	(17,22)	(12,12)	(12,12)	(17,20)	(20,15)
40	(12,15)	(15,12)	(15,12)	(12,15)	(17,12)	(15,20)	(15,17)
41	(10,15)	(15,12)	(12,15)	(12,12)	(15,15)	(15,20)	(15,15)
42	(15,15)	(15,12)	(12,12)	(12,15)	(12,15)	(15,15)	(15,12)
43	(12,15)	(12,10)	(15,10)	(12,15)	(12,15)	(12,15)	(12,12)
44	(10,15)	(7,10)	(10,10)	(10,12)	(12,15)	(12,15)	(10,12)
45	(7,15)	(7,10)	(12,10)	(10,10)	(12,22)	(12,15)	(12,12)
46	(7,15)	(7,10)	(10,10)	(10,15)	(12,12)	(12,15)	(10,10)
47	(7,15)	(7,10)	(15,10)	(12,15)	(12,12)	(12,15)	(10,10)
48	(12,15)	(7,10)	(12,7)	(10,7)	(12,12)	(12,12)	(10,10)

Appendix B

Learning Heating Preferences

The two learning-based thermostat designs introduced in Chapter 4 were based on prior work on home heating by [Shann and Seuken \(2013\)](#). Here we review how they modelled users' heating preferences given real-time prices. When a user wants to heat her house, she needs to decide the temperature that her home will be heated and how much money she is willing to pay for that heating. The *value function* $v(T_t^{int})$ quantifies in pence a user's comfort for a certain temperature T_t^{int} at time t , and the *cost function* $c(T_t^{int}, p_t, T_t^{ext})$ measures how much it will cost to heat to T_t^{int} , given the current price of energy p_t and external temperature T_t^{ext} . The user's *utility* per time period Δt is given by the difference between value and cost multiplied by Δt :

$$u(T_t^{int}, p_t, T_t^{ext}) = (v(T_t^{int}) - c(T_t^{int}, p_t, T_t^{ext}))\Delta t. \quad (\text{B.1})$$

The value function is modelled as a quadratic loss function that has the following parametric form:

$$v(T_t^{int}) = a - b(T^* - T_t^{int})^2, \quad (\text{B.2})$$

where the parameter a is the user's willingness to pay for his most preferred temperature, which is denoted as T^* . The second term, $b(T^* - T_t^{int})^2$ is a loss function that quantifies how much the user suffers from deviations from his most preferred temperature T^* . The parameter b measures the user's *sensitivity* to these temperature deviations.

The cost function is given by the following equation:

$$c(T_t^{int}, p_t, T_t^{ext}) = \lambda p_t r_h (T_t^{int} - T_t^{ext}). \quad (\text{B.3})$$

Here, r_h is heater's power, and λ is the leakage rate of the house. The value of λ captures how well the house is insulated. The cost function approximates the true cost of heating by measuring how much it would cost to keep the temperature at T_t^{int} at time interval t .

Together, this results in the following utility function:

$$u(T_t^{int}, p_t, T_t^{ext}) = (a - b(T^* - T_t^{int})^2 - \lambda p_t r_h(T_t^{int} - T_t^{ext})) \Delta t. \quad (\text{B.4})$$

Using this utility function, it is possible to derive a user's optimal temperature for a given price p_t by taking the derivative with respect to T_t^{int} and solving for T_t^{int} :

$$T_t^{opt}(p_t) = T^* + mp_t, \quad (\text{B.5})$$

where $m = -\frac{\lambda}{2b}$. This means, the optimal temperature T_t^{opt} as a function of the current price p_t is a *decreasing straight line* whose y-intercept is T^* and whose slope is $m = -\frac{\lambda}{2b}$.

At every time step t , the smart thermostat computes the estimated optimal temperature for the current price p_t according to its model of the user's preferences by using the estimates of the most preferred temperature \hat{T}^* and the sensitivity \hat{b} :

$$\hat{T}_{opt}(p_t) = \hat{T}^* - \frac{\lambda}{2\hat{b}} p_t, \quad (\text{B.6})$$

Then, it heats the house to this currently optimal temperature.

References

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.

Alam, M., Gerdin, E. H., Rogers, A., and Ramchurn, S. D. (2015). A scalable, decentralised multi-issue negotiation protocol for energy exchange. In Yang, Q. and Wooldridge, M., editors, *International Joint Conference on Artificial Intelligence (IJCAI)*, pages 1098–1104. AAAI Press.

Alam, M., Ramchurn, S. D., and Rogers, A. (2013). Cooperative energy exchange for the efficient use of energy and resources in remote communities. In *Proceedings of the 2013 International Conference on Autonomous Agents and Multi-agent Systems*, AAMAS ’13, pages 731–738, Richland, SC. International Foundation for Autonomous Agents and Multiagent Systems.

Allcott, H. (2011). Rethinking real-time electricity pricing. *Resource and Energy Economics*, 33(4):820–842.

Allen, J., Quinn, C., and Horvitz, E. (1999). Mixed-initiative interaction. *IEEE Intelligent Systems and their Applications*, 14(5):14–23.

Azaria, A., Gal, Y., Kraus, S., and Goldman, C. V. (2015). Strategic advice provision in repeated human-agent interactions. *Autonomous Agents and Multi-Agent Systems*, 30(1):4–29.

Baker, O. F., Khalil, I., and Kotsis, G. (2009). Agents for energy efficiency in ubiquitous environments. In *Proceedings of the 11th International Conference on Information Integration and Web-based Applications & Services*, pages 674–677. ACM.

Bakker, V., Bosman, M., Molderink, A., Hurink, J., and Smit, G. (2010). Demand side load management using a three step optimization methodology. In *Smart Grid Communications (SmartGridComm), 2010 First IEEE International Conference on*, pages 431–436. IEEE.

Banerjee, N., Rollins, S., and Moran, K. (2011). Automating energy management in green homes. In *Proceedings of the 2Nd ACM SIGCOMM Workshop on Home Networks*, HomeNets ’11, pages 19–24, New York, NY, USA. ACM.

Bang, M., Gustafsson, A., and Katzeff, C. (2007). Promoting new patterns in household energy consumption with pervasive learning games. In *Persuasive Technology*, pages 55–63. Springer.

Bar-Noy, A., Feng, Y., Johnson, M. P., and Liu, O. (2008). When to reap and when to sow—lowering peak usage with realistic batteries. In *Experimental Algorithms*, pages 194–207. Springer.

Barbose, G. L., Goldman, C. A., and Neenan, B. (2004). A survey of utility experience with real time pricing. page 127.

Barker, S., Mishra, A., Irwin, D., Shenoy, P., and Albrecht, J. (2012). Smartcap: flattening peak electricity demand in smart homes. In *Pervasive Computing and Communications (PerCom), 2012 IEEE International Conference on*, pages 67–75. IEEE.

Barkhuus, L. and Dey, A. (2003). Is context-aware computing taking control away from the user? three levels of interactivity examined. In *UbiComp 2003: Ubiquitous Computing*, pages 149–156. Springer.

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.

Berger, C. R. and Calabrese, R. J. (1975). Some explorations in initial interaction and beyond: Toward a developmental theory of interpersonal communication. *Human communication research*, 1(2):99–112.

Bittle, R. G., Valesano, R., and Thaler, G. (1979). The effects of daily cost feedback on residential electricity consumption. *Behavior Modification*, 3(2):187–202.

Bourgeois, J., van der Linden, J., Kortuem, G., Price, B. A., and Rimmer, C. (2014). Conversations with my washing machine: an in-the-wild study of demand shifting with self-generated energy. In *Proceedings of the 2014 ACM International Joint Conference on Pervasive and Ubiquitous Computing*, UbiComp '14, pages 459–470, New York, NY, USA. ACM.

Bradshaw, J., Feltovich, P., and Johnson, M. (2011). *Human-Agent Interaction, In the handbook of human-machine interaction edited by Guy Boy*. Ashgate.

Bradshaw, J. M., Beaument, P., Breedy, M. R., Bunch, L., Drakunov, S. V., Feltovich, P. J., Hoffman, R. R., Jeffers, R., Johnson, M., Kulkarni, S., et al. (2004). Making agents acceptable to people. In *Intelligent Technologies for Information Analysis*, pages 361–406. Springer.

Braun, V. and Clarke, V. (2006). Using thematic analysis in psychology. *Qualitative Research in Psychology*, 3(2):77–101.

Christoffersen, K. and Woods, D. D. (2002). How to make automated systems team players. *Advances in human performance and cognitive engineering research*, 2:1–12.

Clancey, W. (2004). Roles for agent assistants in field science: understanding personal projects and collaboration. *IEEE Transactions on Systems, Man, and Cybernetics, Part C: Applications and Reviews*, 34(2):125–137.

Clear, A., Friday, A., Hazas, M., and Lord, C. (2014). Catch my drift?: achieving comfort more sustainably in conventionally heated buildings. In *Proceedings of the 2014 Conference on Designing Interactive Systems*, DIS '14, pages 1015–1024, New York, NY, USA. ACM.

Clear, A. K., Morley, J., Hazas, M., Friday, A., and Bates, O. (2013). Understanding adaptive thermal comfort: new directions for ubicomp. In *Proceedings of the 2013 ACM International Joint Conference on Pervasive and Ubiquitous Computing, UbiComp '13*, pages 113–122, New York, NY, USA. ACM.

Consolvo, S., McDonald, D. W., Toscos, T., Chen, M. Y., Froehlich, J., Harrison, B., Klasnja, P., LaMarca, A., LeGrand, L., Libby, R., Smith, I., and Landay, J. A. (2008). Activity sensing in the wild: a field trial of ubifit garden. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '08, pages 1797–1806, New York, NY, USA. ACM.

Cook, S. W. (1979). Energy conservation effects of continuous in-home feedback in all-electric homes. *Journal of Environmental Systems*, 9(2):169–173.

Costanza, E., Fischer, J. E., Colley, J. A., Rodden, T., Ramchurn, S. D., and Jennings, N. R. (2014). Doing the laundry with agents: a field trial of a future smart energy system in the home. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '14, pages 813–822, New York, NY, USA. ACM.

Costanza, E., Ramchurn, S. D., and Jennings, N. R. (2012). Understanding domestic energy consumption through interactive visualisation: a field study. In *Proceedings of the 2012 ACM Conference on Ubiquitous Computing*, UbiComp '12, pages 216–225, New York, NY, USA. ACM.

Costanzo, M., Archer, D., Aronson, E., and Pettigrew, T. (1986). Energy conservation behavior: the difficult path from information to action. *American psychologist*, 41(5):521.

Craig, C. S. and McCann, J. M. (1978). Assessing communication effects on energy conservation. *Journal of consumer research*, pages 82–88.

Darby, S. (2006). *The effectiveness of feedback on energy consumption: A review for DEFRA of the literature on metering, billing and direct displays*, volume 486. Environmental Change Institute, University of Oxford.

De Groot, J. I. and Steg, L. (2007). Value orientations and environmental beliefs in five countries validity of an instrument to measure egoistic, altruistic and biospheric value orientations. *Journal of Cross-Cultural Psychology*, 38(3):318–332.

Dobson, J. K. and Griffin, J. A. (1992). Conservation effect of immediate electricity cost feedback on residential consumption behavior. volume 2.

Dütschke, E. and Paetz, A.-G. (2013). Dynamic electricity pricing - which programs do consumers prefer? *Energy Policy*, 59:226–234.

Egan, C. (1999). How customers interpret and use comparative displays of their home energy use. *Proceedings, European Council for an Energy-Efficient Economy*.

Faruqui, A. and Palmer, J. (2011). Dynamic pricing and its discontents. *Regulation*, 34(2):16–22.

Ferber, J. (1999). *Multi-agent systems: an introduction to distributed artificial intelligence*, volume 1. Addison-Wesley Reading.

Fischer, J. E., Ramchurn, S. D., Osborne, M., Parson, O., Huynh, T. D., Alam, M., Pantidi, N., Moran, S., Bachour, K., Reece, S., Costanza, E., Rodden, T., and Jennings, N. R. (2013). Recommending energy tariffs and load shifting based on smart household usage profiling. In *Proceedings of the 2013 International Conference on Intelligent User Interfaces*, IUI '13, pages 383–394, New York, NY, USA. ACM.

Frankcom, N. (2012). Household energy bills unaffordable in less than three years. Technical report.

Froehlich, J., Findlater, L., and Landay, J. (2010). The design of eco-feedback technology. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '10, pages 1999–2008, New York, NY, USA. ACM.

Gal, Y., Kraus, S., Gelfand, M., Khashan, H., and Salmon, E. (2011). An adaptive agent for negotiating with people in different cultures. *ACM Trans. Intell. Syst. Technol.*, 3(1):8:1–8:24.

Geller, E. S. (1981). Evaluating energy conservation programs: is verbal report enough? *Journal of Consumer Research*, 8(3):331–335.

Goodrich, M. A. and Schultz, A. C. (2007). Human-robot interaction: a survey. *Foundations and trends in human-computer interaction*, 1(3):203–275.

Green, R. C., Wang, L., and Alam, M. (2011). The impact of plug-in hybrid electric vehicles on distribution networks: a review and outlook. *Renewable and Sustainable Energy Reviews*, 15(1):544–553.

Haakana, M., Sillanpää, L., and Talsi, M. (1997). The effect of feedback and focused advice on household energy consumption. In *Proceedings, European Council for an Energy-Efficient Economy*.

Heberlein, T. A. and Warriner, G. K. (1983). The influence of price and attitude on shifting residential electricity consumption from on-to off-peak periods. *Journal of Economic Psychology*, 4(1):107–130.

Hines, J., Hungerford, H., and Tomera, A. (1987). Analysis and synthesis of research on responsible environmental behavior: a meta analysis. *The Journal of Environmental Education*, 18(2):1–8.

Hoc, J.-M. (2000). From human–machine interaction to human–machine cooperation. *Ergonomics*, 43(7):833–843.

Holz, T., Dragone, M., and O’Hare, G. (2009). Where robots and virtual agents meet : a survey of social interaction across milgram’s reality-virtuality continuum. *International Journal of Social Robotics*, 1(1):83–93.

Horvitz, E. (1999). Principles of mixed-initiative user interfaces. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI ’99, pages 159–166, New York, NY, USA. ACM.

Hutton, R. B. and McNeill, D. L. (1981). The value of incentives in stimulating energy conservation. *Journal of Consumer Research*, pages 291–298.

Jennings, N. R., Moreau, L., Nicholson, D., Ramchurn, S., Roberts, S., Rodden, T., and Rogers, A. (2014). Human-agent collectives. *Communications of the ACM*, 57(12):80–88.

Jennings, N. R. and Wooldridge, M. J. (1996). Software agents. *IEE review*, pages 17–20.

Kamar, E. S. (2014). *Reasoning effectively under uncertainty for human-computer teamwork*. PhD thesis.

Keshav, S. and Rosenberg, C. (2010). Direct adaptive control of electricity demand. Technical report, Technical Report CS-2010-17, University of Waterloo.

Klein, G., Woods, D. D., Bradshaw, J. M., Hoffman, R. R., and Feltovich, P. J. (2004). Ten challenges for making automation a team player in joint human-agent activity. *IEEE Intelligent Systems*, (6):91–95.

Lam, A. H.-y., Yuan, Y., and Wang, D. (2014). An occupant-participatory approach for thermal comfort enhancement and energy conservation in buildings. In *Proceedings of the 5th International Conference on Future Energy Systems*, e-Energy ’14, pages 133–143, New York, NY, USA. ACM.

Lewis, M. (1998). Designing for human-agent interaction. *AI Magazine*, 19(2):67–78.

Locke, E. A. and Latham, G. P. (2002). Building a practically useful theory of goal setting and task motivation: a 35-year odyssey. *American psychologist*, 57(9):705.

Lu, J., Sookoor, T., Srinivasan, V., Gao, G., Holben, B., Stankovic, J., Field, E., and Whitehouse, K. (2010). The smart thermostat: using occupancy sensors to save energy in homes. In *Proceedings of the 8th ACM Conference on Embedded Networked Sensor Systems*, SenSys '10, pages 211–224, New York, NY, USA. ACM.

Luyben, P. D. (1982). Prompting thermostat setting behavior public response to a presidential appeal for conservation. *Environment and Behavior*, 14(1):113–128.

MacKay, D. (2009). *Sustainable energy - without the hot air*. UIT Cambridge Ltd., Cambridge.

Maes, P., Shneiderman, B., and Miller, J. (1997). Intelligent software agents vs. user-controlled direct manipulation: a debate. In *CHI '97 Extended Abstracts on Human Factors in Computing Systems*, CHI EA '97, pages 105–106, New York, NY, USA. ACM.

McCalley, L. and Midden, C. J. (2002). Energy conservation through product-integrated feedback: the roles of goal-setting and social orientation. *Journal of economic psychology*, 23(5):589–603.

Mennicken, S., Vermeulen, J., and Huang, E. M. (2014). From today's augmented houses to tomorrow's smart homes: new directions for home automation research. In *Proceedings of the 2014 ACM International Joint Conference on Pervasive and Ubiquitous Computing*, UbiComp '14, pages 105–115, New York, NY, USA. ACM.

Microsoft (2008). Being human: human computer interaction in 2020. Technical report, Microsoft.

Miller, C. A. and Parasuraman, R. (2007). Designing for flexible interaction between humans and automation: delegation interfaces for supervisory control. *Human Factors*, 49(1):57–75.

Monigatti, P., Apperley, M., and Rogers, B. (2010). Power and energy visualization for the micro-management of household electricity consumption. In *Proceedings of the International Conference on Advanced Visual Interfaces*, pages 325–328. ACM.

Mountain, D. (2006). The impact of real-time feedback on residential electricity consumption: the hydro one pilot. In *Mountain Economic Consulting and Associates Inc.*

Murphy, R. R. and Schreckenghost, D. (2013). Survey of metrics for human-robot interaction. In *Human-Robot Interaction (HRI), 2013 8th ACM/IEEE International Conference on*, pages 197–198. IEEE.

Myers, K. and Morley, D. (2003). Directing agents. *Agent Autonomy, H. Hexmoor, C. Castelfranchi, and R. Falcone, eds.*, pages 143–162.

Norman, D. (2002). *The design of everyday things*. Basic Books.

Nourbakhsh, I. R., Sycara, K., Koes, M., Yong, M., Lewis, M., and Burion, S. (2005). Human-robot teaming for search and rescue. *Pervasive Computing, IEEE*, 4(1):72–79.

Nwana, H. (1996). Software agents: an overview. *Knowledge Engineering Review*, 11(3):1–40.

Oldewurtel, F., Parisio, A., Jones, C. N., Morari, M., Gyalistras, D., Gwerder, M., Stauch, V., Lehmann, B., and Wirth, K. (2010). Energy efficient building climate control using stochastic model predictive control and weather predictions. In *American Control Conference*, pages 5100–5105, Baltimore, USA.

Pallak, M. S. and Cummings, W. (1976). Commitment and voluntary energy conservation. *Personality and Social Psychology Bulletin*, 2(1):27–30.

Panagopoulos, A. A., Alam, M., Rogers, A., and Jennings, N. R. (2015). Adaheat: a general adaptive intelligent agent for domestic heating control. In *Proceedings of the 2015 International Conference on Autonomous Agents and Multiagent Systems*, AAMAS '15, pages 1295–1303, Richland, SC. International Foundation for Autonomous Agents and Multiagent Systems.

Peffer, T., Pritoni, M., Meier, A., Aragon, C., and Perry, D. (2011). How people use thermostats in homes: a review. *Building and Environment*, 46(12):2529 – 2541.

Peled, N., Gal, Y. K., and Kraus, S. (2015). A study of computational and human strategies in revelation games. *Autonomous Agents and Multi-Agent Systems*, 29(1):73–97.

Pierce, J., Odom, W., and Blevis, E. (2008). Energy aware dwelling: a critical survey of interaction design for eco-visualizations. In *Proceedings of the 20th Australasian Conference on Computer-Human Interaction: Designing for Habitus and Habitat*, pages 1–8. ACM.

Pierce, J. and Paulos, E. (2010). Materializing energy. In *Proceedings of the 8th ACM Conference on Designing Interactive Systems*, DIS '10, pages 113–122, New York, NY, USA. ACM.

Pierce, J. and Paulos, E. (2012). Beyond energy monitors: interaction, energy, and emerging energy systems. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '12, pages 665–674, New York, NY, USA. ACM.

Ramchurn, S., Vytelingum, P., Rogers, A., and Jennings, N. (2012). Putting the ‘smarts’ into the smart grid: a grand challenge for artificial intelligence. *Commun. ACM*, 55(4):86–97.

Ramchurn, S. D., Osborne, M. A., Parson, O., Rahwan, T., Maleki, S., Reece, S., Huynh, T. D., Alam, M., Fischer, J. E., Rodden, T., Moreau, L., and Roberts, S.

(2013). Agentswitch: towards smart energy tariff selection. In *Proceedings of the 2013 International Conference on Autonomous Agents and Multi-agent Systems*, AAMAS '13, pages 1401–1402, Richland, SC. International Foundation for Autonomous Agents and Multiagent Systems.

Reeves, S. (2012). Envisioning ubiquitous computing. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '12, pages 1573–1582, New York, NY, USA. ACM.

Revell, K. M. and Stanton, N. A. (2014). Case studies of mental models in home heat control: searching for feedback, valve, timer and switch theories. *Applied Ergonomics*, 45(3):363 – 378.

Rich, C., Sidner, C., and Lesh, N. (2001). Dynamic pricing and its discontents. *AI Magazine*, 22(4):15–25.

Rodden, T. A., Fischer, J. E., Pantidi, N., Bachour, K., and Moran, S. (2013). At home with agents: exploring attitudes towards future smart energy infrastructures. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '13, pages 1173–1182, New York, NY, USA. ACM.

Rogers, A., Ramchurn, S., and Jennings, N. (2012). Delivering the smart grid: Challenges for autonomous agents and multi-agent systems research. In *Proceedings of the Twenty-Sixth AAAI Conference on Artificial Intelligence*, pages 2166–2172.

Rogers, Y., Connelly, K., Tedesco, L., Hazlewood, W., Kurtz, A., Hall, R., Hursey, J., and Toscos, T. (2007). Why it's worth the hassle: the value of in-situ studies when designing ubicomp. In *Proc. UbiComp*, pages 336–353, Berlin, Heidelberg. Springer-Verlag.

Rosenfeld, A., Zuckerman, I., Segal-Halevi, E., Drein, O., and Kraus, S. (2015). Negochat-a: a chat-based negotiation agent with bounded rationality. *Autonomous Agents and Multi-Agent Systems*, 30(1):60–81.

Sas, C., Whittaker, S., Dow, S., Forlizzi, J., and Zimmerman, J. (2014). Generating implications for design through design research. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '14, pages 1971–1980, New York, NY, USA. ACM.

Sauppé, A. and Mutlu, B. (2015). The social impact of a robot co-worker in industrial settings. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*, CHI '15, pages 3613–3622, New York, NY, USA. ACM.

Scerri, P., Pynadath, D., and Milind, T. (2003). Towards adjustable autonomy for the real world. *Journal of Artificial Intelligence Research*, 2(50):171–228.

Schülke, A., Bauknecht, J., and Häussler, J. (2010). Power demand shifting with smart consumers: a platform for power grid friendly consumption control strategies. In *Smart Grid Communications (SmartGridComm), 2010 First IEEE International Conference on*, pages 437–442. IEEE.

Schwartz, S. (1977). Normative influence on altruism. In L. Berkowitz (Ed.), *Advances in experimental social psychology*, 10:221–279.

Scott, J., Bernheim Brush, A., Krumm, J., Meyers, B., Hazas, M., Hodges, S., and Villar, N. (2011). Preheat: controlling home heating using occupancy prediction. In *Proceedings of the 13th International Conference on Ubiquitous Computing*, UbiComp '11, pages 281–290, New York, NY, USA. ACM.

Shann, M. and Seuken, S. (2013). An active learning approach to home heating in the smart grid. In *Proceedings of the Twenty-Third International Joint Conference on Artificial Intelligence*, IJCAI '13, pages 2892–2899. AAAI Press.

Shenghua, L. (2010). *Interacting with intelligent agents: key issues in agent-based decision support system design*. PhD thesis, University of Jyvaskyla.

Sierhuis, M., Bradshaw, J., Acquisti, R., Hoof, R., and Jeffers, R. (2003). Human-agent teamwork and adjustable autonomy in practice. In *Proceedings of the Seventh International Symposium on Artificial Intelligence, Robotics and Automation in Space (i-SAIRAS*, pages 243–280, Boston, MA. Springer US.

Siero, F. W., Bakker, A. B., Dekker, G. B., and Van Den Burg, M. T. (1996). Changing organizational energy consumption behaviour through comparative feedback. *Journal of environmental psychology*, 16(3):235–246.

Simm, W., Ferrario, M. A., Friday, A., Newman, P., Forshaw, S., Hazas, M., and Dix, A. (2015). Tiree energy pulse: exploring renewable energy forecasts on the edge of the grid. In *Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems*, CHI '15, pages 1965–1974, New York, NY, USA. ACM.

Slavin, R. E., Wodarski, J. S., and Blackburn, B. L. (1981). A group contingency for electricity conservation in master-metered apartments. *Journal of Applied Behavior Analysis*, 14(3):357–363.

Staats, H., Harland, P., and Wilke, H. A. (2004). Effecting durable change a team approach to improve environmental behavior in the household. *Environment and Behavior*, 36(3):341–367.

Staats, H., Wit, A., and Midden, C. (1996). Communicating the greenhouse effect to the public: evaluation of a mass media campaign from a social dilemma perspective. *Journal of environmental management*, 46(2):189–203.

Steg, L. and Vlek, C. (2009). Encouraging pro-environmental behaviour: an integrative review and research agenda. *Journal of environmental psychology*, 29(3):309–317.

Stern, P. C. (1999). Information, incentives, and proenvironmental consumer behavior. *Journal of Consumer Policy*, 22(4):461–478.

Stern, P. C. (2000). Toward a coherent theory of environmentally significant behavior. *Journal of Social Issues*, 56(3):407–424.

Stout, N., Dennis, A. R., and Wells, T. M. (2014). The buck stops there: the impact of perceived accountability and control on the intention to delegate to software agents. *AIS Transactions on Human-Computer Interaction*, 6(1):1–15.

Strengers, Y. (2011). Designing eco-feedback systems for everyday life. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '11, pages 2135–2144, New York, NY, USA. ACM.

Tambe, M., Bowring, E., Pearce, J., Varakantham, P., Scerri, P., and Pynadath, D. (2008). Electric elves: what went wrong and why. *AI Magazine*, 29(2).

The Department of Energy & Climate Change (2009). A consultation on smart metering for electricity and gas. Technical report.

The Department of Energy & Climate Change (2012). Smart metering implementation programme first annual progress report on the roll-out of smart meters. Technical report.

The Department of Energy & Climate Change (2013). Smart meter roll-out for the domestic and small and medium non-domestic sectors. Technical report.

Traum, D. and Rickel, J. (2002). Embodied agents for multi-party dialogue in immersive virtual worlds. In *Proceedings of the First International Joint Conference on Autonomous Agents and Multiagent Systems: Part 2*, AAMAS '02, pages 766–773, New York, NY, USA. ACM.

Truong, N., McInerney, J., Tran-Thanh, L., Costanza, E., and Ramchurn, S. (2013). Forecasting multi-appliance usage for smart home energy management. In *Proceedings of the Twenty-Third International Joint Conference on Artificial Intelligence*, IJCAI '13, pages 2908–2914. AAAI Press.

Tullio, J., Dey, A. K., Chalecki, J., and Fogarty, J. (2007). How it works: a field study of non-technical users interacting with an intelligent system. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '07, pages 31–40, New York, NY, USA. ACM.

US Department of Energy (2003). Grid 2030: a national vision for electricitys second 100 years. Technical report.

US Department of Energy (2006). Benefits of demand response in electricity markets and recommendations for achieving them. Technical report.

US Department of Energy (2008). The smart grid: an introduction. Technical report.

Van Houwelingen, J. H. and Van Raaij, W. F. (1989). The effect of goal-setting and daily electronic feedback on in-home energy use. *Journal of Consumer Research*, pages 98–105.

Völlink, T. and Meertens, R. (1999). De effectiviteit van elektronische feedback over het energie-en waterverbruik door middel van teletekst bij huishoudens.(the effectiveness of electronic feedback on household energy use and water use by means of text tv). *Sociale psychologie en haar toepassingen*, pages 79–91.

Vytelingum, P., Voice, T., Ramchurn, S. D., Rogers, A., and Jennings, N. R. (2011). Theoretical and practical foundations of large-scale agent-based micro-storage in the smart grid. *J. Artif. Intell. Res. (JAIR)*, 42:765–813.

Wang, T. H. and Katzev, R. D. (1990). Group commitment and resource conservation: two field experiments on promoting recycling. *Journal of Applied Social Psychology*, 20(4):265–275.

Weinschenk, S. (2011). *100 things every designer needs to know about people*. New Riders.

Woods, D. and Sarter, N. (2000). Learning from automation surprises and going sour accidents. In *Cognitive engineering in the aviation domain*, N. Sarter and R. Amalberti, eds., pages 327–353. Lawrence Erlbaum.

Wooldridge, M., Jennings, N. R., et al. (1995). Intelligent agents: theory and practice. *Knowledge engineering review*, 10(2):115–152.

Yang, R. and Newman, M. W. (2013). Learning from a learning thermostat: lessons for intelligent systems for the home. In *Proceedings of the 2013 ACM International Joint Conference on Pervasive and Ubiquitous Computing*, UbiComp '13, pages 93–102, New York, NY, USA. ACM.

Yang, R., Newman, M. W., and Forlizzi, J. (2014). Making sustainability sustainable: challenges in the design of eco-interaction technologies. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '14, pages 823–832, New York, NY, USA. ACM.

Yates, S. M. and Aronson, E. (1983). A social psychological perspective on energy conservation in residential buildings. *American Psychologist*, 38(4):435.

Yu, Z., Jia, L., Murphy-Hoye, M. C., Pratt, A., and Tong, L. (2013). Modeling and stochastic control for home energy management. *IEEE Trans. Smart Grid*, 4(4):2244–2255.

Zimmerman, J., Rivard, K., Hargraves, I., Tomasic, A., and Mohnkern, K. (2009). User-created forms as an effective method of human-agent communication. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems*, CHI '09, pages 1869–1878, New York, NY, USA. ACM.