Characterising temporal aspects of residential electricity consumption using statistical learning

Aven Satre-Meloy

St John's College

University of Oxford

A thesis submitted for the degree of

Doctor of Philosophy

Trinity Term, 2019

Abstract

Achieving ambitious climate change mitigation targets requires a comprehensive transformation of the energy system. As part of this transformation, early, rapid, and full decarbonisation of the electricity sector is essential. Residential buildings contribute substantially to electricity consumption, so pathways for electricity system decarbonisation similarly depend on accelerated change in the residential sector.

Delivering further and faster reductions in emissions from residential buildings requires not only reductions in electricity demand but also more responsive demand. Flexibility in the residential sector can support the supply-side transition to renewable energy while also helping reduce the costs of electricity provision during times of peak demand.

Understanding the potential contributions of residential buildings to ambitious mitigation pathways through demand reduction and flexibility requires more detailed characterisations of demand. In particular, it necessitates stronger evidence on the structural and occupant factors that shape electricity demand and its resulting emissions.

This thesis contributes to these research areas and is guided by three primary research questions. The first focuses on improving methods for characterising residential electricity consumption, the second considers advancing our knowledge of factors influencing consumption, and the third relates to the policy implications of having deeper insight into the temporal aspects of demand. The three research questions are:

- What statistical learning approaches can improve efforts to characterise residential electricity consumption patterns?

- What factors explain electricity consumption in residential buildings at annual, daily, and hourly temporal resolutions?

- How does a more detailed understanding of the temporal aspects of residential electricity consumption inform energy efficiency and flexibility interventions?

This thesis consists of four original research papers, each of which makes methodological, empirical, and policy contributions in answering these research questions. The papers' methodological contributions are the development and application of methods from statistical learning and the sub-fields of regularisation and sparsity. These techniques are the subject of a rich literature in statistics disciplines but are not commonly used in energy research. The papers show the utility of these techniques for analysing complex data sets and addressing challenges inherent to building energy modelling.

The papers make empirical contributions to the literature on determinants of electricity consumption through statistical analyses of factors that influence consumption at different temporal resolutions in households. The variety in types of data considered, such as interval electricity, household-level survey, and time-use data, facilitates a deep investigation of the factors that shape residential electricity consumption.

In this way, the papers contribute to the evidence base on effective intervention policies for achieving further demand reduction and flexibility in the residential sector.

Demand-side solutions are underrepresented in the literature on climate change mitigation, but these solutions have considerable potential to achieve ambitious emissions reduction targets. They therefore require a more thorough assessment. This thesis research presents new analytical methods, data, and evidence to investigate how residential buildings can contribute to accelerated decarbonisation pathways.

Acknowledgements

There are many people I wish to thank for their guidance, mentorship, and support throughout the past four years. First and foremost, I would like to thank my supervisors, Dr. Phil Grünewald and Dr. Nick Eyre. If not for Phil's early enthusiasm when I was a master's student, I may never have decided to pursue a doctorate. His interest in my research and his eagerness for me to continue it under his supervision was both humbling and encouraging. I am grateful for the invaluable advice, practical suggestions, and mentorship he has given me over these years, and I cannot thank him enough for his support. I similarly wish to thank Nick for his wisdom and guidance. Nick's helpful feedback often came when I was most struggling to see how the pieces of this thesis would fit together. He gave me clarity as I began to structure my research, and he continued to provide thoughtful comments and encouragement as it developed. Knowing full well that the quality of supervision can vary considerably, I am thankful that I had such constructive and supportive supervisors.

I would also like to express my sincere appreciation for several sources of financial support that have enabled me to undertake this research. I am deeply thankful to the Rhodes Trust, the Sir Peter Elworthy Grant fund, the University of Oxford's Environmental Change Institute, and St John's College. Each of these provided funding, whether for research and travel grants or general living expenses, that allowed me to undertake this research. I am also again indebted to Phil, who hired me as a research assistant in my final year.

Numerous other colleagues have provided mentorship and guidance during my time at Oxford. I am particularly thankful to John Fresen, Becky Ford, Eoghan McKenna, and Pete Walton for sharing their knowledge and providing insightful suggestions for my work. I am grateful for the opportunity to spend a summer conducting research with the U.S.

Department of Energy and wish to thank David Nemtzow and Jared Langevin for the opportunity and for their mentorship.

The Environmental Change Institute has been my home for the past four years, and I have benefited greatly from feedback on my work from the researchers in the Energy Programme, whom I would like to extend my sincere thanks. Thanks, especially, to Marina Diakonova for tirelessly supporting my work and patiently helping me through some of the most daunting analytical aspects of my research. She was always on hand to answer a tough question, and I am deeply appreciative for that. I am also grateful for the chance to learn and work alongside such bright and fun fellow doctoral students, including Laura Brinker, Sam Hampton, Yingqi Liu, Scott MacDonald, and Miriam Zachau-Walker. Thank you for sharing ideas and suggestions over the occasional lunch, coffee, or beer.

There is no question I have drawn on friends both near and far for support when times got tough, and I am deeply appreciative of their willingness to listen, give feedback, and remind me not to get too caught up in it all. There are too many to thank, and I rest assured that you know who you are. I would, however, like to give a special thanks to Patrick Funk for spending countless hours guiding me through dense statistical textbooks and papers and giving me practical feedback on my work. I am deeply grateful for his willingness to help me learn, and I hope to repay the favour some day.

Finally, to my family, I owe you the greatest debt of gratitude. To Liam, Claire, Rena, and Ryan, thank you for inspiring me, teaching me, helping me, visiting me, and keeping me optimistic about the future. To Mom and Dad, your love and support has made all of this possible, and you both have taught me so much about what's really important in the world. While I have not spent as much time with all of you in the past few years as I would have liked, you have given me hope and encouragement when I needed it most. I am so incredibly lucky to have such a supportive, kind, and loving family, and I will never lose sight of that.

Publications arising from this research

The following is a list of published academic journal and conference papers arising from this research.

Journal Articles

- Satre-Meloy, A. (2019), 'Investigating structural and occupant drivers of annual residential electricity consumption using regularization in regression models', *Energy* **174**, 148–168. `https://doi.org/10.1016/j.energy.2019.01.157`.

- Satre-Meloy, A., Diakonova, M., Grünewald, P. (2019a), 'Daily life and demand: An analysis of intra-day variations in residential electricity consumption with time-use data', *Energy Efficiency*, 1–26. `https://doi.org/10.1007/s12053-019-09791-1`.

- Satre-Meloy, A., Diakonova, M., Grünewald, P. (2020), 'Cluster analysis and prediction of residential peak demand profiles using occupant activity data', *Applied Energy* **260**, 1–20. `https://doi.org/10.1016/j.apenergy.2019.114246`.

- Satre-Meloy, A., Langevin, J. (2019), 'Assessing the time-sensitive impacts of energy efficiency and flexibility in the U.S. building sector', *Environmental Research Letters* **14**, 1–12. `https://doi.org/10.1088/1748-9326/ab512e`.

Conference Papers

- Satre-Meloy, A., Diakonova, M., Grünewald, P. (2018), Daily life and demand: New data on behavioral drivers of residential electricity use patterns, *in* 'Proceedings of the 2018 ACEEE Summer Study on Energy Efficiency in Buildings', American

Council for an Energy Efficient Economy, Pacific Grove, CA, pp. 1–12. `https://aceee.org/files/proceedings/2018/index.html#/paper/event-data/p264`.

- Satre-Meloy, A., Diakonova, M., Grünewald, P. (2019b), What makes you peak? Cluster analysis of household activities and electricity demand, *in* 'Proceedings of the 2019 ECEEE Summer Study on Energy Efficiency in Buildings', European Council for an Energy Efficient Economy, Presqu'île de Giens, France, pp. 745–55. `https://www.eceee.org/library/conference_proceedings/eceee_Summer_Studies/2019/4-monitoring-and-evaluation-for-greater-impact/what-makes-you-peak-cluster-analysis-of-household-activities-and-electricity-demand/`.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

1.1 The challenge of the energy transition

The 2018 Special Report from the Intergovernmental Panel on Climate Change (IPCC) lays out in stark detail the transition needed in the power sector to limit global warming to 1.5°C (IPCC, 2018). The scale of this transition is unprecedented. Pathways that are consistent with limiting global warming to 1.5°C require the global electricity sector to emit almost zero greenhouse gas (GHG) emissions by 2050 (Rogelj et al., 2015). This transformation means coal use in the power sector must fall to zero from its current 30% share while renewable electricity must increase to 70–85% of electricity supply in 2050—up from 26% in 2018, of which hydropower supplied around 16% (REN21, 2019).

A rapid transition in the global power sector is complicated by the likelihood of increased electrification in hard to decarbonise sectors, such as transport, heat, and industry. In Europe, for instance, estimates for increased electricity demand due to electrification of transport and heat are around 60–80% of current demand by 2050 (Moraga et al., 2018). Accordingly, the potential for climate change mitigation in heat and transport sectors is clearly dependent on the power sector generation mix. Early and full decarbonisation of the electricity sector is thus a top priority for the energy transition (Rogelj et al., 2015; Eyre et al., 2018). Given required rates of both electrification and power system decarbonisation, reducing total demand is a critical lever for constraining the scale of the energy transition (Eyre et al., 2018).

The building sector is integral to this challenge. Buildings are responsible for a large share of global electricity consumption and resulting emissions. In both Europe and the U.S., for instance, the building sector accounts for around two-thirds of electricity consumption (eurostat, 2009; U.S. EIA, 2018). Both residential and commercial buildings require solutions for further and faster reductions in electricity demand, and far more ambition is needed to deliver emissions reductions in line with the 1.5°C target (Wachsmuth and Duscha, 2019).

1.2 Demand-side solutions for climate change mitigation

Demand-side solutions are underrepresented in the literature on climate change mitigation. Past reports from the IPCC have not provided a thorough assessment of these solutions, especially of their potential to contribute to ambitious mitigation pathways with fewer environmental risks and at lower costs than supply-side solutions (von Stechow et al., 2016).

Previous IPCC reports have incorporated the demand side through techno-economic scenarios that primarily focus on improved end-use energy efficiency. Comparatively less attention in IPCC assessments has been paid to the type, scale, and implementation of demand-side solutions and the associated changes they necessitate, such as shifts in consumption, lifestyles, behaviour, social norms, and technology choice. Yet studies of these kinds of solutions, especially as they pertain to energy use, are the subject of a broad literature in the social and behavioural sciences (e.g., Frederiks et al., 2015). The absence of this literature from past IPCC reports may be due in part to the fact that comprehensive assessments of the potential for demand-side solutions to mitigate climate change, as well as deliver benefits for global health and well-being, are lacking.

Creutzig et al. (2018, p. 260) acknowledge the need to assess the potential for demand-side solutions and address this significant gap in the IPCC's upcoming sixth assessment report (AR6), as they "call for a synthesis of social science and engineering research ... to understand the demand-side potential for climate change mitigation". They dedicate an entire chapter in AR6 to demand, services, and the social aspects of mitigation, the first time an IPCC assessment report will include such a chapter.

This team of researchers emphasises seven key topics that must be addressed in any comprehensive demand-side assessment framework. As a starting point, this assessment must incorporate a detailed characterisation of patterns of demand in energy, mobility, food, and other key systems. Accordingly, the first question for the assessment is "what norms, values, preferences and structural factors shape energy demand and GHG emissions?" (Creutzig et al., 2018, p. 260). Second, the assessment must explore which policy instruments and measures can reduce demand-side GHG emissions while considering specific energy service and socio-economic contexts along with policy constraints. Third, the assessment should investigate how measures can best be implemented and become a part of everyday practice. Fourth, it should estimate the GHG emissions, costs, and potentials associated with specific technologies or systems of provision. The fifth topic and related research question asks how demand-side solutions affect well-being, and the sixth topic focuses broadly on how the demand side can contribute to limiting global warming. Finally, the seventh topic addresses the synergies and trade-offs between demand-side solutions and sustainable development.

These seven topics outline key questions and priorities for research on demand-side solutions in advance of AR6. Importantly, these research areas must be interdisciplinary, including research from fields as diverse as psychology, sociology, political science, economics, industrial ecology, and technology studies, as well as geographically diverse, in order to more carefully consider the cultural and socio-political contexts that underpin the adoption of demand-side solutions. As will be discussed in subsequent chapters, this is and has been an active area of research, especially in the context of energy demand.

These calls are already starting to drive the inclusion of demand-focused scenarios in the Integrated Assessment Modelling (IAM) literature. The IPCC's *Special Report on Global Warming of 1.5° C* (SR15) includes a scenario characterised by rapid reductions in energy demand due to social, business, and technological innovation (IPCC, 2018). In this scenario, rapid decarbonisation of energy supply is enabled by a reduction in global final energy demand of around 40% compared to current levels by 2050 (Grubler et al., 2018).

The authors of this scenario show how down-sizing the global energy system improves the feasibility of decarbonising energy supply. The scenario features both much more rapid adoption of energy efficiency and faster electrification of end uses. Importantly, this scenario is the only one in SR15 not to rely on deployment of bioenergy with carbon capture and storage, a carbon dioxide removal strategy that is modelled in the other scenarios to account for a comparatively slower energy transition. These and related scenarios (e.g., Mundaca et al., 2018) highlight the need for upscaling demand-side solutions to complement the supply-side energy transition and drive deep decarbonisation pathways toward the 1.5°C target.

1.3 Efficiency and flexibility: two key objectives

Decarbonisation pathways undoubtedly require transformations to energy supply, but rapid reductions in electricity demand are also critical for constraining the scale of this transformation (Rogelj et al., 2015). Most of the scenarios analysing demand-side measures in the building sector highlight the importance of energy efficiency improvements to reduce energy demand in line with a 1.5°C target (Mundaca et al., 2018). Furthermore, recent national-level evidence from the E.U. suggests that reducing energy use is an even stronger contributor to climate change mitigation than is indicated in international IAM scenarios (Wachsmuth and Duscha, 2019).

When practitioners and researchers discuss energy efficiency, they most often focus on improvements to the material building fabric and performance of the building's heating, cooling, and other appliances. Related policies include building codes, retrofit programmes, and appliance standards (Ürge-Vorsatz et al., 2013; Wachsmuth and Duscha, 2019). Increasingly, attention is being given to behavioural solutions to reduce energy demand (Mundaca et al., 2018). Gota et al. (2019) highlight both that optimistic mitigation scenarios require deeper behavioural change and that modelling studies underestimate the potential contributions of behaviour change toward mitigation targets (Dietz et al., 2009). Indeed, beyond considering the potential for building retrofits and efficient appliance adoption, the academic energy efficiency community is more recently concerned with ways to limit demand for energy services rather than just improve the end-use efficiency of those services (Darby and Fawcett, 2018). This concept of 'energy sufficiency' is particularly behavioural in concept, placing focus on individual demand for energy services and the potential for demanding less (Willis and Eyre, 2011).

In addition to energy efficiency, another key demand-side solution for decarbonising electricity is flexibility of demand. The concept of demand flexibility maintains that opposed to being uncontrollable, electricity demand can be shifted in time and space in order to match variable energy supply. Grunewald and Diakonova (2018*b*, p. 59) conceive of flexibility simply as "the potential to change power at a certain rate (Watt/hour)". Flexibility presents an alternative way to balance supply and demand, often at lower costs than investing in increased power generation capacity (National Infrastructure Commission, 2016). In general, there are three main demand-side sources of flexibility: energy storage, interconnectors with neighbouring power systems, and demand-side response (DSR) (Newbery, 2018). This last source is also referred to as demand-side management and is defined by the U.S. Federal Energy Regulatory Commission (FERC) as "changes in electric usage by demand-side resources from their normal consumption patterns in response to changes in the price of electricity over time, or to incentive payments designed to induce lower electricity use..." (FERC, 2019). Similarly, the International Energy Agency (IEA) defines building energy flexibility as "the ability [of a building] to manage its demand and generation according to local climate conditions, user needs and energy network requirements" (Jensen et al., 2017, p. 3). When referring to these concepts, this thesis uses the terms flexibility and DSR.

In a similar way as energy efficiency, flexibility and DSR techniques encompass both those that are based on technology improvements, such as the installation of smart appliances or automated building management systems, as well as those that rely on responses from energy users to reduce or shift demand (Strbac, 2008). The former techniques are primarily the subject of a rich technical economic and engineering literature (Aghaei and Alizadeh, 2013; Dupont et al., 2012), but increasingly, non-technical disciplines are exploring the benefits, challenges, and potential of demand flexibility, focusing on important questions about the extent to which people are willing to participate in DSR programmes and their ability to do so (Darby and McKenna, 2012).

A further energy system challenge that is relevant in the context of efficiency and flexibility is that of peak demand, during which times energy provision is often costly and carbon intensive (Torriti, 2017). This is not always the case and is dependent on local energy system dynamics (e.g., Khan et al., 2018), but in contexts where peak demand is met by 'peaker plants', which are typically older, more inefficient, and, as a consequence, more costly and carbon intensive to operate, reducing peak demand becomes an important objective for decarbonising the electricity system (Krieger et al., 2016). The residential sector is a substantial contributor to peak demand in many countries, so developing solutions for peak demand reduction, whether through energy efficiency or flexibility measures, remains an important goal. This is especially pressing given projected increases in the frequency and magnitude of peak demand events under a warming climate (Auffhammer et al., 2017).

While energy efficiency and flexibility are often discussed as mutually beneficial and useful demand-side solutions, there has been limited assessment of their potential when deployed in tandem. Key questions surrounding the integration of these two strategies remain unanswered. While some research has started to explore these questions (e.g., Smith and Managan, 2012), most of the existing literature has afforded only qualitative descriptions of integration opportunities, and there is a need for consistent frameworks to quantitatively assess the combined impacts of efficiency and flexibility on energy demand, especially at regional and national scales (Jensen et al., 2017).

1.4 Characterising residential electricity consumption: the need for new methods and data

In part due to the need to better understand and inform demand-side strategies for electricity system efficiency and flexibility, there is a large and active literature on characterising patterns of electricity demand in buildings. Much of this literature has focused on the physical and social determinants of consumption under different socio-cultural, geographical, or building sector contexts (Jones et al., 2015). Most previous studies apply conventional statistical methods, such as multiple linear regression, to various socio-demographic, dwelling, appliance, and other occupant-related factors to show how these influence patterns of consumption (Fumo and Rafe Biswas, 2015; Swan and Ugursal, 2009). Limitations to smart meter data access among researchers has meant that most of these analyses consider somewhat lower temporal resolutions (months or years) (e.g., Huebner et al., 2016; Ndiaye and Gabriel, 2011), though increased access to interval data in recent years has made studies of much higher temporal resolutions (days, hours, and sub-hour resolutions) more common (e.g., Kavousian et al., 2013; Richardson et al., 2010).

Despite the breadth of this existing literature, significant gaps remain. While many previous studies have confirmed that several key factors, especially those relating to building size and other aspects of the physical dwelling, can explain a high percentage of variance in electricity consumption patterns, this literature has also demonstrated limited overall explanatory power in statistical models (Gram-Hanssen, 2014). This suggests there is still much that is unknown about the various drivers of electricity consumption. One possible reason for this challenge is the complexity of the residential sector given heterogeneity in building sizes and thermal envelope materials, end uses, and occupant behaviours and preferences, all of which interact to drive electricity use. As such, the number of factors that could potentially affect consumption patterns is almost limitless (Hsu, 2015). Swan and Ugursal (2009) add several other challenges: privacy issues have limited the collection of energy data for individual households, and the costs for detailed sub-metering of specific end uses are prohibitive.

Much of the previous work to uncover determinants of electricity demand has collected data through personal interviews, phone and online surveys, or by accessing large databases of electricity, occupant socio-demographic, and building stock records. For this reason, most of the factors that have consistently shown a statistically significant and sizeable influence on electricity consumption patterns are those related to these types of data, such as the number of occupants, their ages, and the household's income, or the size and type of dwelling (Jones et al., 2015). Comparatively much less is known about how the everyday activities performed by occupants relate to electricity consumption patterns (Walker, 2014).

Beyond the difficulties inherent in designing studies to collect these data, from an analysis perspective, incorporating these more nuanced, occupant-related factors into conventional regression models presents challenges because they introduce specific analytical issues, which will be elaborated on further in subsequent chapters. For both of these reasons, comprehensively characterising electricity demand requires new methods that can handle the rich variation in the built environment and new sources of data to shed light on how and when electricity is consumed in homes.

The need for new methods is driven by increasing availability of detailed household and occupant data alongside high-resolution electricity usage data. These are invaluable for identifying key occupant, dwelling, and other household factors that influence patterns of demand, but handling a large number and wide variety of data types in statistical models presents challenges for research (Hsu, 2015).

Interpretability is one such challenge, as not all methods deliver interpretable results. Recent uptake of complex, non-linear statistical models for predicting building electricity usage may improve the accuracy of forecasts, but there is a trade-off here for delivering meaningful insight into factors that influence consumption, as many of these types of models cannot clearly show how changes in such factors associate with changes in electricity consumption (Shmueli, 2010). This lack of interpretability in non-linear approaches makes it difficult to inform policy on how to reduce or shift demand.

Traditional regression approaches, such as ordinary least squares (OLS), can face a different challenge that makes interpretation of their results similarly difficult. Although these methods are computationally straightforward, depending on the structure of the underlying data, OLS can yield relationships between variables that are inflated in magnitude, have the wrong sign, or change considerably given small changes to the underlying data (Friedman et al., 2001). These issues present a different kind of 'interpretability' challenge wherein the relationships between variables are measurable (e.g., they have coefficients with physical significance) but they cannot be trusted given these aforementioned issues. Furthermore, given that OLS yields coefficients for every variable included in a model, when the set of variables to be included is large, OLS models can prove difficult to interpret simply because they will include many variables that are likely irrelevant and unnecessarily add complexity to the model. In other words, OLS models are not always parsimonious even though this is an objective to strive for in statistical analysis. Methods that can address these multiple forms of the interpretability challenge in the energy domain are thus needed. While such methods exist and are commonplace in statistics disciplines, their application to energy data is rare (Hsu, 2015).

There is also a need for new types of data to shed light on important questions for electricity demand. In particular, the central question of what electricity is used for in homes is one for which limited empirical evidence exists (Grunewald and Diakonova, 2018*b*; Powells et al., 2014). Time-use data show promise for exploring this question, and these data are now more frequently incorporated into models of electricity demand (e.g., Widén and Wäckelgård, 2010), but there remains a notable gap in the literature of empirical studies on how occupants' use of energy services influences the temporal aspects of consumption (Anderson and Torriti, 2018; Grunewald and Diakonova, 2018*a*). One suggested reason for this gap is given by Shove and Walker (2014, p. 42), who argue that much of the energy demand literature is motivated by a school of thought that understands energy use as "either the cause or the consequence of changing political, economic and technical systems". Placing energy production and use in this context causes the important question of 'what is energy used for?' to fall out of focus. Asking this question takes a different view—one that examines energy use not as an outcome of the intersections of technological innovation and political-economic systems but rather as something defined by social practices. This view, and its implications for the research undertaken in this thesis, will be more fully explored in Chapter 2.

Finally, the question of timing is an especially important consideration for energy demand research. Characterising electricity consumption patterns at different temporal resolutions can inform different solutions. The factors that influence consumption across months and years may not be the same that explain daily or hourly variations in demand. Understanding these differences is useful for informing different demand-side solutions. For instance, deeper insight into the factors that influence consumption across years, such as the size and efficiency of a dwelling, can inform strategies to reduce demand in absolute terms over long timescales. These same factors, however, may not be as useful for reducing demand at different times of day, which is necessary to limit peak demand and also to make demand more responsive. The need for new data at higher temporal resolutions is thus similarly important for informing these efforts.

1.5 Thesis research question, aims, and structure

Returning to Creutzig et al.'s (2018) framework for assessing the demand-side potential for climate change mitigation, this thesis makes contributions to several of the key research priorities set forth. First, regarding the authors' point that the starting place for any comprehensive assessment of demand-side solutions is a detailed characterisation of demand, this thesis contributes to that aim for the residential building sector. It focuses specifically on electricity demand given the pressing need to decarbonise the power sector. To contribute to a deeper understanding of the various social and physical factors that shape patterns of residential electricity consumption, this thesis aims to improve the analytical approaches for characterising electricity consumption, identify key factors that shape demand, and provide insight into the policy implications of considering the temporal aspects of residential electricity use. As such, there are three overarching research questions guiding the work undertaken for this thesis, which are as follows:

Methodological research question: What statistical learning approaches can improve efforts to characterise residential electricity consumption patterns?

Empirical research question: What factors explain electricity consumption in residential buildings at annual, daily, and hourly temporal resolutions?

Policy research question: How does a more detailed understanding of the temporal aspects of residential electricity consumption inform energy efficiency and flexibility interventions?

In answering these three research questions, this thesis applies novel methods from the field of statistical learning to characterise residential building electricity consumption patterns at successively resolved timescales, starting with annual then daily then hourly variations in consumption. It shows how recent advances in statistical learning, especially from the field of statistical learning with sparsity, can improve modelling efforts by addressing important analytical challenges that arise in energy demand research. Next, it uses these new methods to identify key structural and occupant-related factors that shape residential electricity demand at these varying temporal resolutions. In doing so, it contributes to Creutzig et. al's second research topic by informing strategies for demand reduction and flexibility, both of which are integral policy measures for reducing emissions from the building sector and supporting a rapid supply-side transformation to a low-carbon power system. Finally, the thesis contributes to the fourth topic mentioned above by estimating the potential cost and GHG emissions reduction potentials for specific energy efficiency and flexibility measures when deployed at a national scale.

The thesis consists of four original research papers. The first three introduce and apply statistical learning methods to explore drivers of electricity consumption in households at the three temporal resolutions described above. These papers build upon one another in terms of the methods and techniques used to characterise electricity demand as well as the data used in the models. The fourth paper provides further commentary on this shift from an annual to hourly temporal resolution by arguing for the need to adopt a time-sensitive framework when quantifying the potential savings from building efficiency and flexibility measures. This policy-focused paper provides evidence on the potential benefits of energy efficiency, peak demand reduction, and demand flexibility in the building sector, and it places its analysis at the scale of the national U.S. building stock. A graphical overview of the thesis research questions and the structure of the four papers, including the methods and data they use, is presented as an infographic in Figure 1.1.

The thesis also includes several chapters that provide background and context on related literature and the design of the thesis research. These are organised as follows: Chapter 2 is a literature review that continues on from Section 1.2 to discuss recent literature on demand-side solutions for climate change mitigation and theoretical frameworks for conceptualising residential electricity consumption before providing an overview of literature relevant to the three thesis research questions. Chapter 3 provides a methodology for the thesis. It first introduces the field of statistical learning and justifies their use in this thesis. It then presents the structure of the thesis and the relationships between the four empirical papers. These are presented next in their published form, though some formatting changes have been applied for consistency. Chapter 8 presents a summary of the core findings and research implications, taken together from the four papers, and discusses limitations for the thesis as a whole and directions for future research. Chapter 9 concludes.

1.6 Novelty of thesis

This thesis makes several novel contributions to the energy demand literature in terms of methodological improvements, empirical findings, and policy implications. The central link between the thesis research papers is the application of advanced statistical learning methods to multivariate energy and household data. While these methods are studied extensively in statistics and econometric disciplines, they are not often used to characterise energy demand or better understand its temporal patterns and the influential factors driving these. I argue that this is a missed opportunity for energy researchers, as the methods are particularly useful for handling numerous analytical challenges that are common in quantitative energy demand research, and they also deliver more interpretable results, both in comparison to non-linear or 'deep learning' methods and conventional regression techniques.

The thesis' empirical contributions also provide novel insight into the factors influencing demand at different temporal resolutions. In particular, the papers contribute to the literature on how occupant time use affects electricity demand in households throughout the day and during evening peak hours, providing detailed evidence on the important question of what energy is used *for*. The relationship between time use and electricity use is one for which little empirical evidence exists (e.g., Widén et al., 2009). Strengthening this evidence base is thus one of the novel contributions of this thesis.

In identifying key factors that explain residential electricity consumption patterns, the papers present novel insights for demand-side policies and strategies, especially those that pertain to targeted interventions to reduce or shift demand. In this way, the papers show how more detailed characterisations of demand can both inform the design and delivery of energy efficiency and flexibility measures in the residential sector while also aiding policymakers in comparing the time-sensitive economic and climate benefits of these types of interventions.

As part of publishing the thesis' four papers, several novel data sets and methodological frameworks are being shared with the energy research community. These include: 1) a database of occupant time use and electricity use, which is the first such database of its kind incorporating both occupant activities and high-resolution electricity data collected in parallel from U.K. households; and 2) a new model and evaluation framework for assessing the benefits of energy savings measures with time-sensitive impacts in the U.S. building sector. Both of these outputs will support other energy researchers undertaking similar studies of electricity demand in the residential sector.

1.7 Chapter summary

This chapter introduced the challenge of the energy transition, the importance of developing demand-side solutions for mitigation, and the central role of energy efficiency and flexibility for decarbonising the electricity system. It discussed how detailed characterisations of residential electricity demand are necessary to identify strategic opportunities for achieving these goals. The chapter then introduced the three primary research questions guiding this research and explained the thesis' aims and structure. The following chapter continues on from this introduction and presents an overview of the literature relevant to the thesis' research questions.

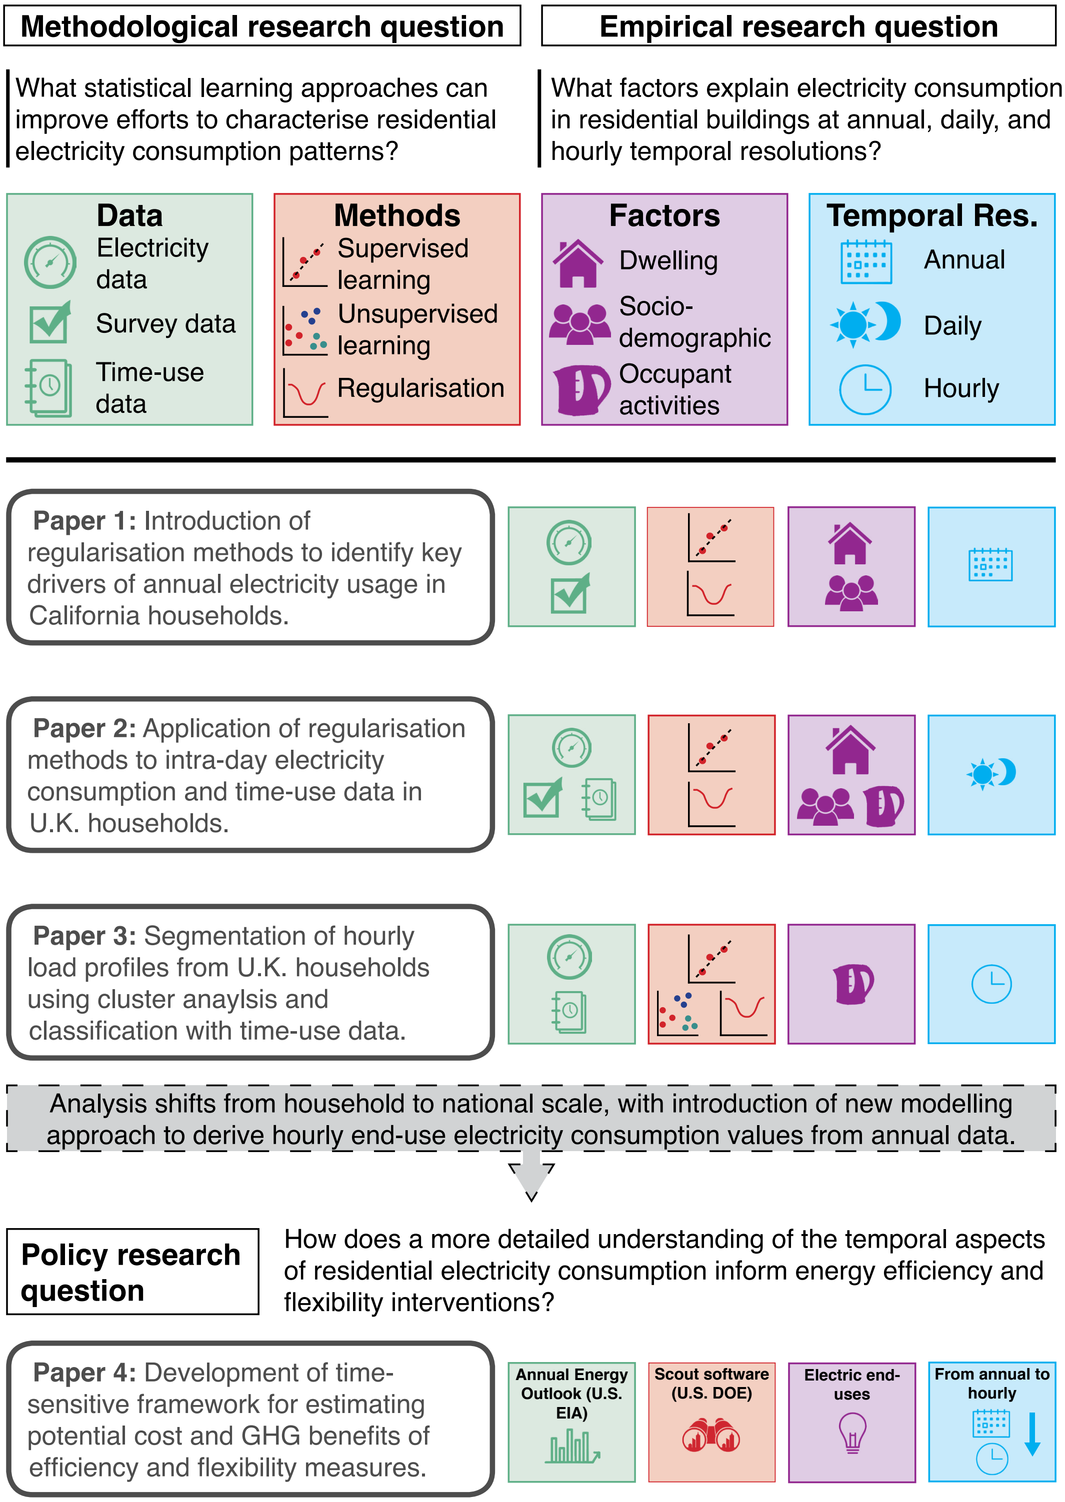

Figure 1.1: Infographic overview of the four empirical papers.

1.8 References

Aghaei, J. and Alizadeh, M.-I. (2013), 'Demand response in smart electricity grids equipped with renewable energy sources: A review', *Renewable and Sustainable Energy Reviews* **18**, 64–72.

Anderson, B. and Torriti, J. (2018), 'Explaining shifts in U.K. electricity demand using time use data from 1974 to 2014', *Energy Policy* **123**, 544–557.

Auffhammer, M., Baylis, P. and Hausman, C. H. (2017), 'Climate change is projected to have severe impacts on the frequency and intensity of peak electricity demand across the United States', *Proceedings of the National Academy of Sciences* **114**(8), 1886–1891.

Creutzig, F., Roy, J., Lamb, W. F., Azevedo, I. M. L., Bruine de Bruin, W., Dalkmann, H., Edelenbosch, O. Y., Geels, F. W., Grubler, A., Hepburn, C., Hertwich, E. G., Khosla, R., Mattauch, L., Minx, J. C., Ramakrishnan, A., Rao, N. D., Steinberger, J. K., Tavoni, M., Ürge-Vorsatz, D. and Weber, E. U. (2018), 'Towards demand-side solutions for mitigating climate change', *Nature Climate Change* **8**(4), 260–263.

Darby, S. and Fawcett, T. (2018), Energy sufficiency: An introduction, Technical report, European Council for an Energy Efficient Economy, Oxford, United Kingdom.

Darby, S. J. and McKenna, E. (2012), 'Social implications of residential demand response in cool temperate climates', *Energy Policy* **49**, 759–769.

Dietz, T., Gardner, G. T., Gilligan, J., Stern, P. C. and Vandenbergh, M. P. (2009), 'Household actions can provide a behavioral wedge to rapidly reduce U.S. carbon emissions', *Proceedings of the National Academy of Sciences* **106**(44), 18452–18456.

Dupont, B., Tant, J. and Belmans, R. (2012), Automated residential demand response based on dynamic pricing, *in* '2012 3rd IEEE PES Innovative Smart Grid Technologies Europe (ISGT Europe)', IEEE, Berlin, Germany, pp. 1–7.

eurostat (2009), *Harmonised European Time Use Surveys: 2008 Guidelines*, Office for Official Publications of the European Communities, Luxembourg. OCLC: 698893637.

Eyre, N., Darby, S. J., Grünewald, P., McKenna, E. and Ford, R. (2018), 'Reaching a 1.5°C target: Socio-technical challenges for a rapid transition to low-carbon electricity systems', *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* **376**(2119).

FERC (2019), 'Reports on demand response and advanced metering', `https://www.ferc.gov/industries/electric/indus-act/demand-response/dem-res-adv-metering.asp`.

Frederiks, E. R., Stenner, K. and Hobman, E. V. (2015), 'Household energy use: Applying behavioural economics to understand consumer decision-making and behaviour', *Renewable and Sustainable Energy Reviews* **41**, 1385–1394.

Friedman, J., Hastie, T. and Tibshirani, R. (2001), *The Elements of Statistical Learning*, Vol. 1 of *Springer Series in Statistics*, Springer, Berlin, Germany.

Fumo, N. and Rafe Biswas, M. (2015), 'Regression analysis for prediction of residential energy consumption', *Renewable and Sustainable Energy Reviews* **47**, 332–343.

Gota, S., Huizenga, C., Peet, K., Medimorec, N. and Bakker, S. (2019), 'Decarbonising transport to achieve Paris Agreement targets', *Energy Efficiency* **12**(2), 363–386.

Gram-Hanssen, K. (2014), 'New needs for better understanding of household's energy consumption - Behaviour, lifestyle or practices?', *Architectural Engineering and Design Management* **10**(1-2), 91–107.

Grubler, A., Wilson, C., Bento, N., Boza-Kiss, B., Krey, V., McCollum, D. L., Rao, N. D., Riahi, K., Rogelj, J., De Stercke, S., Cullen, J., Frank, S., Fricko, O., Guo, F., Gidden, M., Havlík, P., Huppmann, D., Kiesewetter, G., Rafaj, P., Schoepp, W. and Valin, H. (2018), 'A low energy demand scenario for meeting the 1.5°C target and sustainable development goals without negative emission technologies', *Nature Energy* **3**(6), 515–527.

Grunewald, P. and Diakonova, M. (2018*a*), 'The electricity footprint of household activities–implications for demand models', *Energy and Buildings* **174**, 635–641.

Grunewald, P. and Diakonova, M. (2018*b*), 'Flexibility, dynamism and diversity in energy supply and demand: A critical review', *Energy Research & Social Science* **38**, 58–66.

Hsu, D. (2015), ‘Identifying key variables and interactions in statistical models of building energy consumption using regularization’, *Energy* **83**, 144–155.

Huebner, G., Shipworth, D., Hamilton, I., Chalabi, Z. and Oreszczyn, T. (2016), ‘Understanding electricity consumption: A comparative contribution of building factors, socio-demographics, appliances, behaviours and attitudes’, *Applied Energy* **177**, 692–702.

IPCC (2018), Summary for Policymakers, *in* V. Masson-Delmotte, P. Zhai, H.-O. Pörtner, D. Roberts, J. Skea, P. Shukla, A. Pirani, W. Moufouma-Okia, C. Péan, R. Pidcock, S. Connors, J. Matthews, Y. Chen, X. Zhou, M. Gomis, E. Lonnoy, Maycock, M. Tignor and T. Waterfield, eds, ‘Global Warming of 1.5°C. An IPCC Special Report on the Impacts of Global Warming of 1.5°C above Pre-Industrial Levels and Related Global Greenhouse Gas Emission Pathways, in the Context of Strengthening the Global Response to the Threat of Climate Change, Sustainable Development, and Efforts to Eradicate Poverty’, World Meteorological Organization, Geneva, Switzerland.

Jensen, S. Ø., Madsen, H., Lopes, R., Junker, R. G., Aelenei, D., Li, R., Metzger, S., Lindberg, K. B., Marszal, A. J., Kummert, M., Bayles, B., Mlecnik, E., Lollini, R. and Pasut, W. (2017), *Annex 67: Energy Flexible Buildings*, International Energy Agency, Paris, France.

Jones, R. V., Fuertes, A. and Lomas, K. J. (2015), ‘The socio-economic, dwelling and appliance related factors affecting electricity consumption in domestic buildings’, *Renewable and Sustainable Energy Reviews* **43**, 901–917.

Kavousian, A., Rajagopal, R. and Fischer, M. (2013), ‘Determinants of residential electricity consumption: Using smart meter data to examine the effect of climate, building characteristics, appliance stock, and occupants’ behavior’, *Energy* **55**, 184–194.

Khan, I., Jack, M. W. and Stephenson, J. (2018), ‘Analysis of greenhouse gas emissions in electricity systems using time-varying carbon intensity’, *Journal of Cleaner Production* **184**, 1091–1101.

Krieger, E. M., Casey, J. A. and Shonkoff, S. B. C. (2016), ‘A framework for siting and dispatch of emerging energy resources to realize environmental and health benefits: Case study on peaker power plant displacement’, *Energy Policy* **96**, 302–313.

Moraga, J. L., Le Coq, C., Mulder, M. and Schwenen, S. (2018), Gas and the electrification of heating & transport: Scenarios for 2050, Technical report, Centre on Regulation in Europe, Brussels, Belgium.

Mundaca, L., Ürge-Vorsatz, D. and Wilson, C. (2018), 'Demand-side approaches for limiting global warming to 1.5°C', *Energy Efficiency* .

National Infrastructure Commission (2016), Smart Power, Technical report, London, United Kingdom.

Ndiaye, D. and Gabriel, K. (2011), 'Principal component analysis of the electricity consumption in residential dwellings', *Energy and Buildings* **43**(2-3), 446–453.

Newbery, D. (2018), 'Shifting demand and supply over time and space to manage intermittent generation: The economics of electrical storage', *Energy Policy* **113**, 711–720.

Powells, G., Bulkeley, H., Bell, S. and Judson, E. (2014), 'Peak electricity demand and the flexibility of everyday life', *Geoforum* **55**, 43–52.

REN21 (2019), Renewables 2019 Global Status Report, Technical report, REN21 Secretariat, Paris, France.

Richardson, I., Thomson, M., Infield, D. and Clifford, C. (2010), 'Domestic electricity use: A high-resolution energy demand model', *Energy and Buildings* **42**(10), 1878–1887.

Rogelj, J., Luderer, G., Pietzcker, R. C., Kriegler, E., Schaeffer, M., Krey, V. and Riahi, K. (2015), 'Energy system transformations for limiting end-of-century warming to below 1.5°C', *Nature Climate Change* **5**(6), 519–527.

Shmueli, G. (2010), 'To explain or to predict?', *Statistical Science* **25**(3), 289–310.

Shove, E. and Walker, G. (2014), 'What Is energy for? Social practice and energy demand', *Theory, Culture & Society* **31**(5), 41–58.

Smith, K. and Managan, K. (2012), Energy efficiency and demand response: Working together in an integrated approach to managing energy, Technical report, Institute for Building Efficiency, Johnson Controls, Washington, D.C.

Strbac, G. (2008), 'Demand side management: Benefits and challenges', *Energy Policy* **36**(12), 4419–4426.

Swan, L. G. and Ugursal, V. I. (2009), 'Modeling of end-use energy consumption in the residential sector: A review of modeling techniques', *Renewable and Sustainable Energy Reviews* **13**(8), 1819–1835.

Torriti, J. (2017), 'The risk of residential peak electricity demand: A comparison of five European countries', *Energies* **10**(3), 385.

Ürge-Vorsatz, D., Petrichenko, K., Staniec, M. and Eom, J. (2013), 'Energy use in buildings in a long-term perspective', *Current Opinion in Environmental Sustainability* **5**(2), 141–151.

U.S. EIA (2018), Annual Energy Outlook 2018 with projections to 2050, Technical report, U.S. Energy Information Administration, Washington, D.C.

von Stechow, C., Minx, J. C., Riahi, K., Jewell, J., McCollum, D. L., Callaghan, M. W., Bertram, C., Luderer, G. and Baiocchi, G. (2016), '2°C and SDGs: United they stand, divided they fall?', *Environmental Research Letters* **11**(3), 1–15.

Wachsmuth, J. and Duscha, V. (2019), 'Achievability of the Paris targets in the E.U. - The role of demand-side-driven mitigation in different types of scenarios', *Energy Efficiency* **12**(2), 403–421.

Walker, G. (2014), 'The dynamics of energy demand: Change, rhythm and synchronicity', *Energy Research & Social Science* **1**, 49–55.

Widén, J., Lundh, M., Vassileva, I., Dahlquist, E., Ellegård, K. and Wäckelgård, E. (2009), 'Constructing load profiles for household electricity and hot water from time-use data—Modelling approach and validation', *Energy and Buildings* **41**(7), 753–768.

Widén, J. and Wäckelgård, E. (2010), 'A high-resolution stochastic model of domestic activity patterns and electricity demand', *Applied Energy* **87**(6), 1880–1892.

Willis, R. and Eyre, N. (2011), *Demanding Less: Why We Need a New Politics of Energy*, Park Lane Press, London, United Kingdom.

Chapter 2

Literature Review

2.1 Introduction

This chapter provides a review of literature relevant to the research undertaken for the thesis. Because the empirical papers each have detailed literature reviews pertinent to their individual research questions, this chapter will focus more broadly on the literature that leads to the three research questions introduced in the previous chapter. Throughout this chapter, however, I indicate which of the thesis' papers include more detailed reviews of specific literature so that the reader is aware of where further background information is presented.

This chapter proceeds as follows: Section 2.2 continues on from the introduction to provide a more comprehensive review of recent work on characterising and quantifying demand-side solutions for climate change mitigation. Section 2.3 then provides background on theoretical approaches for conceptualising household energy consumption. Next, Section 2.4 reviews literature relevant to the methodological thesis research question on different approaches for characterising residential electricity demand. It includes a discussion of the overall taxonomy of energy models, further detail on statistical modelling approaches, and a review of key sources of data used in different models. Section 2.5 then reviews literature relevant to the empirical thesis research question regarding key factors that influence consumption. This section focuses on the temporal resolution of previous studies, introduces the types of explanatory factors considered, and highlights findings from decades of research on residential electricity consumption. Finally, Section 2.6 provides a review of key policy considerations for improving the efficiency and flexibility of residential electricity consumption. This section discusses literature relevant to the third thesis research question, examining benefits and challenges of energy efficiency and flexibility while also reviewing literature on interventions to change demand.

2.2 Demand-side solutions for climate change mitigation

The literature on pathways to achieve ambitious mitigation targets has, until recently, focused primarily on supply-side technology change and carbon dioxide removal options. Increasingly, however, the academic and policy communities have begun to assess the potential for demand-side approaches to drive deep decarbonisation pathways in line with the 1.5°C target. Evidence for this shift is apparent in the recent publication of and call for special issues in leading journals focused on demand-side solutions. The journal *Energy Efficiency* in 2019 published a collection of papers in the issue "Demand-side policies and socio-technical pathways to limit global warming below 1.5°C". Similarly, the journal *Environmental Research Letters* has opened a call for submissions for a special issue to be published in 2020 entitled, "Focus on demand-side solutions for transitioning to low-carbon societies". The following section will summarise this emerging field of literature, focusing specifically on these new special issues and related research efforts underway in both the U.K. and U.S. These areas have been the subject of much previous research, however, and the evidence base underpinning these is reviewed in more depth in subsequent sections and in the thesis' four papers.

The purpose of the special issue published in *Energy Efficiency* is "to broaden and strengthen the evidence base on the role of demand-side policies, measures, and corresponding mitigation pathways to limit global warming to 1.5°C" (Mundaca et al., 2018, p. 344). Among the research questions that are addressed through the various papers included in the issue are several that are specific to building sector decarbonisation. These include: "How can 1.5°C pathways be pursued through (radical) changes in household consumption?" and "What policy options will encourage deep decarbonization pathways in the building sector?".

The special issue's introductory editorial summarises the key insights from papers across three themes: policy interventions, demand-side measures, and methodological approaches. Highlights of those papers that are particularly relevant to residential building energy consumption are provided here.

Several papers in the special issue highlight the importance of stringency in demand-side policies to meet ambitious mitigation targets. Among the policies focused on specifically are building codes, performance standards, and behavioural policies. Brown and Li (2019) show how stringent building codes can lead to much higher thermal efficiencies in single-family homes and apartments, while Méjean et al. (2019) demonstrate the importance of building and energy efficiency standards for reducing energy demand in residential buildings. Wachsmuth and Duscha (2019) show that codes can enable rapid reductions in building sector emissions if new buildings and retrofits are required to meet the most stringent thermal-efficiency standards available. Performance standards for home appliances are shown to be similarly effective for mitigation. Sonnenschein et al. (2019) find that modest standards for several key residential appliances are complimentary to carbon pricing policies and could achieve the same emissions reductions as existing carbon prices.

The papers also make clear that behaviour-focused policy interventions are necessary alongside technical change for delivering ambitious mitigation targets. For instance, Sonnenschein et al. (2019) argue that behavioural aspects of consumption should be addressed along with performance standards and carbon pricing, and the authors suggest labelling programs and green defaults as examples of policy options that can address behavioural anomalies, such as limited attention spans among consumers. In the context of residential heating, Knobloch et al. (2019) suggest that policy instruments depend on assumptions related to behavioural decision-making, and the authors argue that failing to account for behavioural nuances can lead to misinformed policy recommendations. Moberg et al. (2019) stress the importance of more regulatory-based approaches for policies that target consumer behaviour, arguing that existing policies that rely on self-governance and nudging are insufficient for encouraging further behaviour change.

Regarding sectoral decarbonisation pathways for the buildings sector, there is consistent agreement among papers that rapid improvements in energy efficiency are needed to limit global warming to 1.5°C. Wachsmuth and Duscha (2019) show via an index decomposition that the ambitious mitigation pathways for the building sector in the E.U. are in part due to lifestyle changes such as a slower increase in housing sizes but are primarily attributable to higher energy efficiency levels and rates of electrification. Strict standards for new buildings and retrofits also contribute to deep decarbonisation scenarios in the E.U. Wilson et al. (2018) identify disruptive innovations that can deliver different strategies for building decarbonisation, such as device interconnectivity to optimise electricity usage, improved thermal performance through the adoption of smart heating controls, and reduced demand for materials given the emergence of the sharing economy and device convergence (which refers to the trend of devices being used for increasingly many purposes, thus reducing the need for multiple devices; the example of smart phones is given).

These kinds of disruptive innovations feature in the so-called 'Low Energy Demand (LED)' scenario, which projects a 40% reduction in global final energy demand by 2050 in line with a 1.5°C target (Grubler et al., 2018). For the residential building sector, the authors show thermal energy demand in the developed world of around 160 MJ/m^2 by 2050, down from around 600 MJ/m^2 in 2020. This reduction assumes doubling the current thermal retrofit rate from 1.4% to 3% per year. Regarding appliances and lighting, the scenario projects electricity demand for these is reduced by 20% in the developed world, which is due both to reductions in energy service demand from the disruptive innovations mentioned above, as well as reductions in end-use energy intensity from the dissemination of high-efficiency lighting and appliances (see Supplementary Notes 3–4 in Grubler et al. (2018) for more detail). Finally, both Brown and Li (2019) and Knobloch et al. (2019) show that stringent building codes, high-thermal insulation, and highly efficient appliances, especially air conditioners and heat pumps, feature prominently in deep decarbonisation pathways for the building sector and require rapid deployment in order to keep building-sector emissions in line with a 1.5°C target.

A final theme explored in the papers in the special issue is that of methodological approaches for analysing and evaluating demand-side policies and related mitigation pathways. Four categories of methodological approaches are discussed: system modelling, decomposition analyses, bottom-up approaches, and stakeholder approaches. The third category of bottom-up approaches is the most relevant for the research undertaken in this thesis. These approaches are defined by Mundaca et al. (2018, p. 357) as "disaggregated characterizations and analyses (including quantitative modeling) of end-use sectors, services, or technologies that aim to provide policy-relevant knowledge". The types of bottom-up approaches demonstrated in the papers include those that estimate emissions reduction potentials for disruptive low-carbon innovations (Wilson et al., 2018), forecast demand and emissions reduction targets for the U.S. electricity sector using the U.S. National Energy Modelling System (NEMS) (Brown and Li, 2019), and simulate technology diffusion for residential heating systems up to 2050 (Knobloch et al., 2019). The variety of approaches taken in these papers highlights that many types of methods are necessary for fully evaluating demand-side solutions and providing the necessary policy insight to enact them.

While it has yet to be published, the call for a special issue in *Environmental Research Letters* similarly highlights the need for a stronger evidence base on demand-side solutions for ambitious mitigation pathways. As expressed by the guest editors for the special issue, the motivation for a more comprehensive assessment of demand-side solutions for limiting global warming below 1.5°C is threefold. First, modelling scenarios lack short-term mitigation options at the individual, community, and societal scale, especially those that represent demand-side behavioural solutions. Second, these scenarios do not sufficiently incorporate solutions that stem from disciplines other than those from the conventional techno-economic paradigm; important findings from the social sciences, especially psychology, sociology, and anthropology, are not reflected in mitigation assessments. Third, there is a need for more evidence on what solutions work well and under what conditions in order for the IPCC to understand and assess demand-side solutions more adequately. The call aims to bring together evidence through systematic reviews of literature in these areas in order to create a knowledge base that can aid the systems modelling community in better representing demand-side options in mitigation scenarios.

In addition to these special issues and the publications contained therein, other institutions and research agencies are taking on the question of energy demand and its role in the low-carbon transition. The newly-established Centre for Research into Energy Demand Solutions (CREDS) in the U.K. recently released its first major publication, which addresses the question of how changes in energy demand can contribute to a net-zero carbon energy system in the U.K. (Eyre and Killip, 2019). The report examines how energy demand plays a role in decarbonisation efforts in buildings, transport, industry, and electricity sectors, and it sets forth recommendations for policy. The broad summary recommendations from the report include: 1) prioritisation of energy demand solutions alongside those dealing with energy supply; 2) greater consideration and promotion of the health and economic benefits of demand-side solutions; 3) further scaling up of policies that prove effective in reducing emissions; 4) development of longer-term plans for innovation in energy demand; 5) increased activity in demand-side policy delivery across different levels of government; and 6) involvement of a wide range of stakeholders to strengthen support for energy demand solutions.

In the U.S., similar efforts are underway. For example, the U.S. Department of Energy (DOE) is developing a strategy for advancing the role of buildings in supporting a low-carbon electricity grid. DOE recently released a comprehensive report on Grid-interactive Efficient Buildings (GEB), which provides an overview of the opportunities and benefits of optimising the integration of distributed energy resources to meet both occupant and electric grid needs (Neukomm et al., 2019). The report includes an assessment of the different grid services that can be provided by more flexible building end uses in order to position the building sector as increasingly central to a low-carbon energy transition. At a time when much of the federal policy support for the supply-side transition to renewable energy has stalled in the U.S., demand-side solutions have received increased attention given their potential benefits in terms of cost savings, grid reliability and security, and meeting the diverse needs and preferences of occupants. While DOE's GEB strategy is one that is primarily driven by technology development rather than social or behavioural change, it is still a clear example of an emerging focus on demand-side solutions in U.S. energy policy.

Across these literature streams, there is support for much more detailed assessments of the potential contributions demand-side solutions can make toward ambitious mitigation targets. These are seen not as secondary but as equally important to supply-side transitions in order to keep the 1.5°C target within reach. The authors of these various papers and reports make clear that demand-side solutions can enable deeper and faster decarbonisation while reducing its associated costs. Demand-side measures are abundant, and supporting policies already exist to deploy these, but they are not usually incorporated in systems models and their mitigation scenarios. An additional challenge referred to specifically in Mundaca et al. (2018) is the issue of integrating and assessing behavioural or lifestyle changes in modelling studies. The authors argue that existing evidence from social and behavioural sciences makes clear that typical conceptions of behaviour from rational choice theory are wholly insufficient for deeply understanding the dynamics of energy use. Further support for this argument is given in Creutzig et al. (2016). Incorporating new insights on behaviour into energy models is thus increasingly important for informing and assessing demand-side solutions.

2.3 Theoretical frameworks for conceptualising household energy consumption

While there is growing recognition among techno-economic disciplines that a deeper understanding and more accurate representation of human behaviour is critical for achieving ambitious climate change mitigation targets, energy social scientists and theorists have long held that other disciplinary approaches are key for understanding energy consumption and energy-related behaviour (Sovacool, 2014). As such, the literature on theoretical approaches to conceptualising residential energy consumption is vast. One of they key questions examined in this literature, which is wide-ranging and explores many different avenues for understanding human-energy dynamics, is where agency is rooted in the energy system (Gram-Hanssen, 2014). This question is an important one, as different ways of thinking about agency—in other words, who or what has the capacity to influence or change something—can lead to contrasting conclusions about the most effective ways to reduce or shift household energy demand. Understanding the behavioural, lifestyle, and cultural factors underpinning household energy consumption is essential for informing energy efficiency and flexibility programmes and policies, and the social sciences have a central role to play in this effort.

This section thus provides a brief review of several theoretical approaches to understanding household energy consumption, focusing on how different approaches locate agency in the energy system. What follows is not an exhaustive review but rather an attempt to highlight some of the emerging views that contextualise residential energy consumption and hold insights for how to change it. The section first introduces some of the more conventional approaches for thinking about agents in the energy system. Next, it introduces practice theory, which diverges from past approaches and has recently received considerable attention from energy social scientists given its usefulness for understanding energy consumption in everyday life (Gram-Hanssen, 2014; Røpke, 2009; Shove and Walker, 2010).

2.3.1 Locating agency in the energy system

Much of the literature on conceptualising residential energy use can be categorised into three different perspectives: technological, individual, and socio-cultural. Each of these is reviewed in turn.

Approaches that consider technology as the agent of change in the energy system are commonly applied in technical or engineering disciplines. These approaches believe that energy system improvements, including both those for individuals (e.g., comfort, cleanliness, convenience) and those for the system (e.g., automation, lower emissions, resilience) are defined and solved with technology. Technological approaches argue that the best tools for making these improvements to the energy system are technologies, such as those that make the system more efficient, smart, automated, low-carbon, and so on. Importantly, this approach considers users who attempt to control the system as 'problems' (Higginson et al., 2014).

A critique of the technology perspective is given in Janda (2011), who argues that most of the solutions for reducing energy use are purely architectural, relying on stakeholders such as architects and engineers but excluding building owners, operators, developers, and users. This last group of stakeholders is argued to play an essential role in the built environment, but the users—individual people—are often overlooked.

Understanding the role of energy users is thus central to the individual or behavioural perspective, which focuses on the individual as the agent in the energy system. Two common interpretations of how this agency manifests itself consider the individual, or agent, as either a rational economic actor (Blasch et al., 2019) or a psychologically and behaviourally consistent actor (Brown and Macey, 1983). In the first of these interpretations, individuals have access to all the essential information needed to make economically rational decisions. In the second approach, the focus is on individuals' attitudes, norms, intentions, and knowledge and assumes that their behaviour will be psychologically consistent with these values. In other words, it assumes that people behave as they intend to (Ajzen and Fishbein, 1980; Desmedt et al., 2009).

These models of consumer choice and decision, which focus on the individual as the change agent, suggest that wider societal or systemic change results from individual consumers acting in response to behaviour change stimuli such as pricing, information, education, and social norms. This approach has been applied in many studies on energy consumption that use incentives or tailored information and feedback to reduce energy use (Abrahamse et al., 2007; Abrahamse and Steg, 2009; Pothitou et al., 2017). This literature includes studies that are exploratory and seek to better understand the links between energy-related behaviour and individuals' values, attitudes, knowledge, habits, and socio-demographics (e.g., Huebner et al., 2016) and those that are experimental or include interventions to test individuals' responses to information or incentives (Dolan and Metcalfe, 2015). Literature specific to both of these study types is reviewed in greater detail in Section 2.6.

The third perspective or approach locates agency at the wider societal or cultural level within which individuals make choices that affect energy use. These approaches are variously referred to as lifestyle (Gram-Hanssen, 2014), contextual (Hargreaves, 2011), cultural (Gram-Hanssen, 2010*a*), or structural (Shove, 2010) approaches. These ways of conceptualising energy use suggest that society retains more deterministic agency than individual consumers and that energy consumption is structured by economic and cultural capital, suggesting that change can be induced by focusing on these structures, either at the macro level using legislation and policymaking tools, at the social level, where marketing is particularly effective, or at the local level, where community-based or social networking approaches are most appropriate.

Socio-cultural approaches focus on consumption as a way of showing identity. While there are indeed links between energy consumption and displays of social status in society, energy consumption is often a much more routinised and mundane form of consumption, so theoretical perspectives that focus too much on 'conspicuous' consumption, which many claim cultural approaches do, fail to reveal much insight about more 'inconspicuous' consumption activities, such as those that consume energy in the home (Groncow and Warde, 2001). For this reason, the approach is critiqued for having limits in both its application to energy studies as well as its ability to yield interventions that radically change household energy consumption (Barr et al., 2005).

Increasingly, energy researchers acknowledge that the more routinised and structured parts of consumer activities must be more in focus to understand household energy consumption (Gram-Hanssen, 2010*b*; Hargreaves, 2011; Shove and Walker, 2014; Wilhite et al., 2000). This acknowledgement has precipitated the establishment of a practice theoretical framework to better interrogate residential energy consumption by recognising that agency is shared across all three levels—between individual consumers, the technologies they use, and the broader social context for their consumption.

2.3.2 Practice theory and its application to energy demand

Practice theory focuses on "the collective structures of practices and on what guides the practices people perform in their everyday lives" (Gram-Hanssen, 2014, p. 94). Agency is understood to be embedded within social practices. In the context of energy use, the consumption of energy itself is not a practice; instead, the many different activities that people do in their homes that consume energy, such as cooking, washing, and recreating, are the focus of analysis. There is not one common understanding of what practice theory is, and there are many different contributions to this framework from various disciplines, but this review of practice theory draws on work from Gram-Hanssen (2009, 2014) and Shove and Walker (2014), who further develop the ideas discussed in Schatzki (1996, 2002) and Reckwitz (2002).

As defined by Schatzki, "a practice is a temporally evolving, open-ended set of doings and sayings linked by practical understandings, rules, teleoaffective structures, and general understandings..." (Schatzki, 2002, p. 87). Gram-Hanssen (2009, p. 47) explains that in defining practices, Schatzki "introduces a hierarchy with doings and sayings at the basic level, and with collections of sayings and doings forming a level of tasks, which, in turn, at a higher level with several tasks, can form projects." Both tasks and projects can be composed of different sets of sayings and doings, and a participant in a practice can carry out actions at all three of these levels. Practices can be occasional or routine doings and saying but also include rare sayings and doings, tasks, and projects.

Consider a simple household practice such as cooking. Sayings and doings (for instance, talking with family members about what to cook), tasks (purchasing groceries), and projects (additional tasks necessary for cooking, such as preparing ingredients and getting a recipe) are all part of the cooking 'practice' but comprise different levels of this hierarchy and are in some cases, such as with purchasing groceries, less routine than preparing ingredients or having conversations about what to cook. Even more rare, the task of purchasing new cooking equipment happens much more infrequently than preparing ingredients, but these sets are all part of the same practice. Gram-Hanssen (2009) notes that when performing practices, participants engage in social activities with those who they interact with directly, such as other family members in the case of cooking, but also with all other people performing a certain practice (for example, cooking practices are often shared within cultural contexts).

Schatzki's definition of practice suggests four links that hold together sayings and doings: practical understandings, rules, teleoaffective structures, and general understandings. Here again, as with the definition of practice theory, there is no complete agreement about the elements of a practice. Different theorists have varying names and explanations for these elements. Practical understanding is linked to practical intelligibility, which refers to the fact that individuals are guided often not by what is the rational or the normatively correct thing to do but instead by what makes sense for the individual. Practical understanding is about knowing what to do or how to react to something—for example, the know-how in regulating indoor temperature by adjusting a thermostat or opening the windows. Rules are explicit rules for how to do things, not implicit rules that might influence individual behaviours. An example of rules in the context of energy use might be written rules for connecting an electric or gas meter in a new home to set up a utility service (Gram-Hanssen, 2009). Teleoaffective structures refers to how practices are often goal-oriented or directed by normative beliefs or opinions. Norms hold together practices, as in the case of norms of cleanliness and washing practices, but individuals engaging in practices do not have to be aware of the teleoaffective structures of these to take part. In the example of washing practices, individuals may not refer to norms of cleanliness when describing washing, but nonetheless their actions are often guided by norms. General understandings refers to more widely-held beliefs, values, attitudes, or concerns, such as those that are established more broadly by religion or culture.

Practice theory maintains that all of these elements are equally important for understanding practices and that they are dual in nature, meaning they should be seen as structures that sustain practices at the same time as they are sustained and developed by those engaging in the practices (Gram-Hanssen, 2009). Understanding how these elements influence practices and are influenced by practitioners can be useful for determining the most effective ways to change practices that use energy. Where educational or information campaigns may help shift individuals' knowledge about their energy-consuming practices, more subtle interventions might rely on the powers of social norms for changing energy-consuming behaviours. Equally, interventions should also take into consideration embodied habits and the role of technologies in energy consumption because these likely play a prominent role in structuring energy consumption at the household level.

Locating agency for household energy consumption with energy practices themselves rather than with technologies, individuals, or at the level of lifestyle or culture can provide insights for how to change energy consuming behaviours when it is clear that groups with similar values, attitudes, and lifestyles, as well as similar household characteristics and technologies, use widely varying amounts of energy. Indeed, many studies have shown that this is often the case (Gram-Hanssen, 2010*a*; Santin et al., 2009; Steemers and Yun, 2009). This is likely why practice theory has been taken as a relevant and insightful theoretical framework from which to understand and change household energy consumption.

Returning to the previous chapter's argument that there is a need for more detailed, empirical evidence on *what energy is used for* (Shove and Walker, 2014; Walker, 2014), practice theory emerges as a useful framework for thinking about household energy use and puts this important question in focus. This thesis does not directly apply a practice theoretical approach or methodology, but because of its principal focus on energy use behaviour and the activities (or practices) that underpin it, practice theory is a useful framework for understanding and conceptualising energy use in the various empirical studies that are presented as part of this thesis. These points will be returned to in the review of energy efficiency and flexibility policy in Section 2.6.

2.4 Methodological approaches for characterising residential electricity consumption

This section reviews literature relevant to the thesis' methodological research question on how novel statistical learning techniques can improve efforts to characterise electricity consumption patterns in the residential sector. Again, large portions of the literature relevant to this question are reviewed in detail in the empirical papers, so this section primarily provides an introduction to the literature and indicates in which papers further reviews are given. One important caveat to the following review is that it does not include discussions of some alternative methodological approaches, especially those that are qualitative in nature. A brief discussion of these approaches is provided in the following chapter (Section 3.4).

2.4.1 The taxonomy of modelling approaches

Numerous quantitative modelling approaches have been used to characterise energy consumption in residential buildings (Foucquier et al., 2013). While all building energy models aim to quantify energy use as a function of input variables or parameters, the approaches vary widely depending on the scale of the analysis and the sources of data used in the models. An example of scale is whether the model aims to determine national energy supply requirements or to estimate changes in consumption for an individual building given its characteristics. Sources of data are similarly wide-ranging. These can be large-scale surveys or simulations of electricity consumption and related factors or specific, physical measurements of individual buildings and detailed socio-demographic data from their occupants.

Given these considerations, energy modelling approaches can generally be grouped into two broad categories: top-down or bottom-up. Swan and Ugursal (2009) present a detailed review of these, and the first paper in this thesis does so as well, but given that modelling is the central methodological approach in this thesis, some framing material is also presented here. A schema of the taxonomy of modelling approaches for residential electricity demand is presented in Figure 2.1.

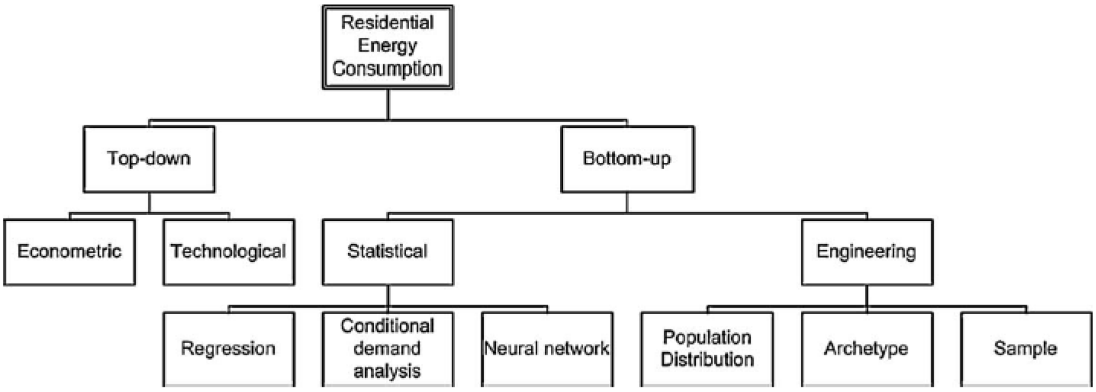

Figure 2.1: Taxonomy of residential energy modelling approaches from Swan and Ugursal (2009). An adapted version of this figure appears in paper 1.

Top-down models are typically used to attribute energy consumption (for instance, in the residential sector) to characteristics of the entire building stock. These are used to forecast long-term changes in consumption primarily to inform supply-side considerations, such as the need to build additional power plants to meet demand. Variables included in these models are mostly macroeconomic indicators, such as regional- or national-level gross domestic product (GDP), climate projections, or new construction rates. Top-down models can be further categorised as econometric or technological models, which are differentiated mainly in terms of their inputs, where econometric models use price and income parameters and technological models attribute consumption to technology trends, such as rates and types of appliance ownership.

The main strength of top-down models is that they are relatively easy to construct and have fewer data requirements because they are based on macroeconomic indicators, which are often more readily-available than other data. They are also useful for long-term forecasting and thus are primarily used for supply analysis. The central weakness of these models is the coarse nature of their analytical approach, which does not explicitly consider specific end uses and cannot account for more complex, individual building interactions that might be relevant to other energy research questions.

Bottom-up models differ from top-down ones by modelling at the scale of individual homes and building end uses. These approaches typically develop a model for a sample of homes and then extrapolate to larger geographic areas. As with top-down models, there are two main categories of bottom-up models: statistical and engineering. Statistical approaches establish relationships between energy consumption and model inputs mathematically, whereas engineering approaches rely on thermodynamic relationships to establish the energy consumption of different end uses (for instance, based on power ratings and equipment use). Both types of bottom-up approaches require more detailed input data than do top-down models, but an advantage of this detail is the ability to represent more complex and nuanced relationships between inputs and energy consumption, especially those related to occupant behaviour.

In bottom-up statistical models it is straightforward to determine the typical contribution of different end uses to energy consumption patterns. For example, detailed electric billing data can be modelled with household survey data on appliance ownership to determine which types of end uses are most able to explain variations in electricity consumption for an individual household. The coefficients that are solved for in statistical models can be interpreted directly because they indicate how much of a change in consumption is expected given a change in ownership or use of specific appliances. The drawback of statistical methods, however, is that they usually require large sample sizes, access to billing or otherwise metered consumption data, and they suffer some specific statistical challenges, which are further discussed in the next chapter as well as in paper 1.

Bottom-up engineering methods, on the other hand, explicitly calculate or simulate building energy consumption and do not rely on historical consumption data (though these data can be used for model calibration). In this way, they are highly accurate for individual dwellings, and this makes them the best option for evaluating the impact of new technologies. The weaknesses of these methods are that they require very detailed data inputs, can be computationally intensive, and are difficult to scale to large samples. To represent occupant behaviour, such as the switching on and off of individual appliances, these models typically use stochastic programming, such as Monte Carlo methods or the Markov-chain technique, to produce outputs that are matched to probability distributions found in the real world (McKenna and Thomson, 2016). Determining these probabilities is a step in the model development process. A related approach that can be applied in the context of energy demand is agent-based modelling, in which individual actors and their behaviours are represented as software agents that simulate aspects of real behaviour, which in turn can be based on individual decision-making rules that are determined via economic optimisation models or are based on ethnographic research on energy consumption behaviour (Hiteva et al., 2018). The application of engineering models to energy demand is wide-ranging and has received sustained attention from the energy research community. For a detailed review, interested readers should see Crawley et al. (2008) and Nguyen et al. (2014).

The specific methods shown under statistical and engineering approaches in Figure 2.1 have different properties and applications, and a detailed review of these is beyond the scope of this chapter. The methods applied in this thesis research, however, are primarily bottom-up statistical regression techniques. A brief review of these is presented next.

2.4.2 Statistical modelling of electricity demand

Among the different statistical modelling approaches, regression techniques are the most common. They are useful for characterising building energy consumption because they are relatively easy to implement, they require less computational power than more complex approaches, they have satisfactory predictive ability, and they benefit from increased availability of both electricity and accompanying household and building data (Fumo and Rafe Biswas, 2015). As mentioned in the previous section, regression analysis is an approach that seeks to find a functional relationship (represented by an equation) between a response or dependent variable and predictor or independent variables. Linear regression is the simplest and most commonly-used type of regression, but there are numerous extensions to the linear regression setting to handle situations like multiple response variables, a qualitative rather than quantitative response variable, or a non-linear relationship between predictors and the response. The first three empirical papers present more detailed reviews of these extensions.

The two other statistical methods shown in Figure 2.1 are conditional demand analysis and artificial neural networks, and these are both briefly reviewed in the first paper. These different methods all fall under the broad umbrella of statistical learning, which is introduced in Chapter 3. A related statistical method called clustering is yet another common statistical learning approach for characterising patterns of electricity demand, especially of high-resolution electric load data. A detailed review of literature applying this approach is given in the third paper.

These methods are all similar, however, in their ability to handle many different sources of data and provide much more detailed insight into the relationships between electricity demand and structural and occupant factors, especially those related to occupant preferences, behaviours, and activities. In the next section, an introduction to some of these types of factors and how they have been incorporated in previous studies is given.

2.4.3 Key sources of data used for characterising residential electricity consumption

The specific sources of data used in different modelling approaches varies considerably, and availability of or access to different sources of data constrains the choice of modelling approach. Several key types of data used for residential electricity modelling include the following:

Electricity consumption data

There are several different ways of accessing and collecting electricity consumption data, each with advantages and disadvantages. The first way is collecting data from detailed monitoring of individual dwellings (Kavousian et al., 2013; McLoughlin et al., 2012; Ndiaye and Gabriel, 2011). The advantage of this approach is more accurate and reliable estimates for electricity usage, but the disadvantage is the challenge of accessing and collecting these data, which requires intrusive methods and can only be done at a small scale. This is especially the case when individual sub-metering of plug loads is desired to access granular, appliance-level electricity consumption data. The second method is collecting electricity data from an energy supplier (Bartiaux and Gram-Hanssen, 2005; Bartusch et al., 2012; Wyatt, 2013) or from a household's billing data (Sanquist et al., 2012). The main benefit of these approaches is that they give researchers access to much larger samples of data to analyse, but the drawback is that these data are sometimes unreliable, either because of faulty meter readings or poor estimates from utilities. In the empirical papers included in this thesis, data are collected both through monitoring of individual dwellings as well as from energy suppliers. Further discussion of the advantages and disadvantages of these methods is provided in the papers.

Household survey data

Household surveys are used to gather detailed information on dwelling and occupant characteristics, such as the physical aspects and thermal efficiency of the dwelling, the appliance stock, and occupant-specific information, including detailed appliance usage data, energy behaviours, and occupant socio-demographic data. In many cases, researchers access these data from existing national-scale databases (Brounen et al., 2012; Hsu, 2015; Huebner et al., 2016; McLoughlin et al., 2012). In other studies, researchers design their own survey instruments for collecting these data (Cayla et al., 2011; Wallis et al., 2016). As was the case with electricity consumption data, there are similarly several different ways to collect survey data when the instrument is designed specifically for the study. Personal interviews (Gram-Hanssen et al., 2004; Tso and Yau, 2007), phone surveys (Ndiaye and Gabriel, 2011), and questionnaires, delivered either in the mail or online, (Baker and Rylatt, 2008; Bartusch et al., 2012; Kavousian et al., 2013) are all used to collect survey data. In the first three empirical papers in this thesis, household survey data are collected via original survey instruments, which are delivered online. The fourth paper accesses data from national databases. More detail on the different data collection approaches for household survey data is included in the papers.

Time-use data

Time-use data refer to detailed records of what individuals do or how they spend their time in the normal course of daily life. These data are less-commonly associated with modelling residential electricity consumption, but in the past decade, energy researchers have begun to incorporate time-use data into models of electricity demand to understand how the timing of home occupancy and occupant activities affects electricity consumption patterns (Torriti, 2014). These data can be useful, as well, for understanding how long-term changes in how people spend their time influence electricity consumption (De Lauretis et al., 2017; Jalas and Juntunen, 2015), or for validating other high-resolution energy demand models (McKenna and Thomson, 2016). The application of time-use data for characterising residential electricity consumption in the thesis' papers is one of the novel aspects of this research. The second and third empirical papers thus provide more thorough reviews of the literature on time-use modelling of electricity consumption.

2.5 Empirical findings on factors influencing residential electricity consumption

This section reviews literature relevant to the thesis' empirical research question on factors that influence residential electricity consumption at annual, daily, and hourly temporal resolutions. As with the previous section, much of this literature is reviewed in the empirical papers. This section thus briefly reviews two broad themes that are relevant for the empirical research question. First, it discusses the importance of temporal resolution for influencing the findings and implications of energy demand research. Next, it reviews the broad categories of factors included in previous studies and highlights key evidence on how these affect consumption.

2.5.1 The importance of temporal resolution

The temporal resolution of electricity demand research is an important consideration because it influences the type of modelling approach used, the specific inputs or factors that are included, and the findings and implications of the research. For example, it is important to consider whether the analysis will focus on explaining variations in a household's baseload demand or its discretionary usage. Baseload is made up of continuous loads and standby appliances (Nelson, 2008), whereas discretionary usage refers to electricity consumption occurring due to active residential behaviour, such as appliance use, lighting, and space heating or cooling (Jin et al., 2016). As will be discussed in detail in papers 2 and 3, there are important methodological considerations for how to pre-process electricity usage data depending on whether the focus of the analysis is explaining changes in baseload versus discretionary usage.

Similarly, the implications for research focused on longer-term electricity patterns will differ from research examining variations over shorter time periods. As was discussed in Chapter 1, while understanding what factors explain usage patterns over months and years may be helpful for identifying levers for further reductions in use over these timescales, the same knowledge may not be helpful for informing pathways for flexibility in demand over hours and days.

Historically, much of the literature on determinants of residential electricity consumption has focused on longer temporal resolutions, such as annual or monthly patterns of electricity use, primarily due to data availability. Evidence for this comes from Jones et al. (2015), who present the most recent comprehensive review of studies investigating the effects of socio-demographic, dwelling, and appliance factors on electricity consumption. Of the 37 international studies they include in their review, 29 of these specifically focus on factors influencing monthly, seasonal, or annual variations in electricity consumption. Of these, 27 include socio-demographic factors, 23 include dwelling factors, and 18 include appliance factors. While the inclusion of different types of factors is clearly also dependent on data availability, this review suggests that previous studies have mostly been concerned with exploring links between socio-demographic or dwelling factors and annual or monthly electricity usage patterns.

In recent years, however, increased access to smart meter data among researchers has supported a shift in energy research toward analyses of daily and hourly electricity use. This point is supported by findings from a literature review undertaken for the third paper in this thesis, which summarises the literature on daily and hourly load profile segmentation across 30 studies (see Table 6.1). Of the papers reviewed, all but four were published after 2010. While the literature search was not systematic but was instead based on a snowball approach, this suggests a growing interest and ability among energy researchers to collect and analyse data at higher temporal resolutions.

Access to higher resolution data has opened up new avenues for research to investigate how socio-demographic and dwelling factors are associated with hourly and daily patterns of electricity use. Most of the existing evidence for the links between these factors, such as socio-economic status and family structure or dwelling size and type, and electricity use comes from models with lower temporal resolutions (Kavousian et al., 2013). An important avenue of research is understanding how these affect consumption over the course of the day. Paper 2 examines this question in detail.

To a similar extent, the type of occupant behavioural factors included in energy demand models is influenced by the temporal resolution of the analysis. In some cases, behavioural factors can have long-term impacts. Purchasing energy efficient appliances or engaging in habitual energy-saving practices, such as regularly turning down the thermostat or turning off the lights, can affect electricity usage over the course of months and years in addition to influencing daily patterns of use (Gardner and Stern, 2009; Steemers and Yun, 2009). Occupant behavioural factors that aim to measure actual use of electricity, via appliances or other electricity-consuming activities, have similarly been incorporated into electricity models at varying temporal resolutions, though time-use data, for example, are more often incorporated in higher-resolution models (McKenna and Thomson, 2016; Torriti, 2014). As occupant behaviour is a key focus of this thesis, the first three papers present more detailed discussions of how the type of occupant behavioural factors included in models is influenced by their temporal resolution.

Finally, as paper 4 argues, model temporal resolution has implications for estimating the value associated with reductions or shifts in consumption. Because the cost and emissions associated with electricity use vary temporally (on both daily and seasonal timescales), so do the potential savings for measures that reduce electricity use at different times of days or in different seasons (Mims et al., 2017). For this reason, the temporal resolution chosen for a model that aims to quantify the potential value of different energy efficiency or flexibility measures is of central importance. Paper 4 demonstrates this empirically and discusses the advantages of adopting a time-sensitive modelling framework in studies of the technical potential of energy efficiency and flexibility.

2.5.2 Categories of factors and their effects on consumption

As was mentioned in Section 2.4.3, a wide range of household data has been incorporated in statistical models of energy consumption. A non-exhaustive list of the categories of factors examined in previous studies includes physical dwelling factors, appliance ownership and usage factors, occupant socio-demographic characteristics, aspects related to occupant attitudes, literacy, values, and preferences in the context of energy use, and detailed aspects of occupant behaviour. This section presents a brief review of specific factors included in previous research and evidence for their effects on electricity consumption. It should be noted, however, that while this section mostly reviews several meta-analyses of factors influencing electricity use, more detailed reviews of specific studies and their findings are given in the papers.

Physical dwelling factors

Factors relating to the physical characteristics of a dwelling are frequently included in electricity consumption research. Those that have been studied the most (in terms of number of citations in which they appear) are: dwelling floor area, type (e.g., detached, semi-detached), and age; number of rooms, bedrooms, and floors; factors related to the dwelling's thermal envelope; and presence or absence of various heating and cooling equipment (Jones et al., 2015).

Dwelling size and type are the two of these that show the most unequivocal effect on electricity consumption. In Jones et al.'s (2015) review, 22 studies include floor area, and 19 of these show it as having a significant, positive effect on residential electricity consumption, meaning it is associated with higher usage. Three studies show it having no effect. Dwelling type (for increasing levels of detachment) shows a significant, positive effect in 12 studies, with zero citing a negative or no effect on consumption. Dwelling age has more mixed evidence, whereas the number of rooms, bedrooms, and floors also generally show positive effects (because these are typically used as proxies for the size of a dwelling when floor area is not available). The presence of space and water heating equipment as well as air conditioning are factors that all show a significant, positive effect in most studies in which they are included.

In terms of the explanatory or predictive power of this category of factors, most studies show that these can explain a comparatively high percentage of variance when compared to the other categories (Fumo and Rafe Biswas, 2015; Gram-Hanssen, 2014; Huebner et al., 2016).

Appliance factors

Appliances are included as factors in two main forms: ownership and use. Many appliances, such as refrigerators, consume electricity irrespective of their use, which is why the total number of appliances owned is expected to influence electricity consumption. Some studies measure appliance ownership as a single total, whereas others specify categories of appliances, and yet others only include specific appliances. Appliance use, on the other hand, is often measured in minutes per day or hours per week (e.g., TVs, computers, other IT equipment) or number of times used per week (e.g., number of showers, dishwasher or washing machine cycles per week) (Bedir et al., 2013).

Jones et al. (2015) show that appliance ownership is more often included in studies than appliance use, perhaps owing to the fact that data on appliance ownership are more easily obtained and more reliable than actual or estimated usage data. Regarding specific appliances, the most frequently studied include desktop computers, TVs, electric ovens, refrigerators/freezers, tumble dryers, and washing machines. There is also substantial evidence that ownership of these shows a significant, positive effect on electricity consumption. In fact, Jones et al. (2015) find that in none of the studies reviewed do any of these show a significant, negative effect on consumption (but in several cases they find that these show no effect). While there is less evidence on appliance use, there is consistency across studies that it is associated with higher electricity consumption. This is especially true for tumble dryer, washing machine, and dishwasher use.

Several recent studies in the U.K. have taken an alternate approach to measure the effects of appliance use on electricity consumption and confirm these results (Palmer, Terry and Kane, 2013, Palmer et al., 2013). These used detailed, circuit-level monitoring of individual household appliances to understand patterns of electricity use at high temporal resolutions. The objective of these studies was to understand patterns of appliance usage and the resulting impact on total electricity consumption and peak demand. Findings from this research indicate that cold appliances, audiovisual equipment, lighting, cooking appliances, and washing appliances are the largest end use contributors to annual electricity consumption in U.K. households, whereas cooking appliances, lighting, and audiovisual devices contribute the most to U.K. peak demand.

As is the case with dwelling factors, previous studies have found appliance factors can explain a comparatively high percentage of variance in residential electricity use (Bedir et al., 2013; Carlson et al., 2013; Huebner et al., 2016).

Socio-demographic factors

Socio-demographic factors, also referred to as socio-economic factors, include the following: number of occupants, family structure (e.g., number of children, adults, and elderly people), age, education level, and employment status of the occupants (or of one designated 'responsible' member of the household), household income, and tenure type. Of these, the number of occupants and household income are by far the most frequently studied factors (included in 23 and 21 studies, respectively, out of the 37 reviewed in Jones et al. (2015)).

These two factors are also those which most consistently show a significant, positive effect on electricity consumption in the literature. Occupant age also consistently shows a positive effect. Presence of children, education level, and tenure show more mixed effects on variations in consumption.

While some socio-demographic factors show strong and statistically significant associations to electricity use, as a category of factors, these have been shown to explain lower percentages of variance than dwelling and appliance factors in previous studies (Chen et al., 2013; Huebner et al., 2016; Ndiaye and Gabriel, 2011).

Occupant attitude factors

Occupant attitudes, norms, and preferences are not commonly included in quantitative studies of residential electricity demand. For instance, this was not a category of factors included in Jones et al.'s (2015) review, likely owing to the fact that collecting data on these is more challenging than for other categories and because there is more subjectivity in the data collection process. Studies that have included these have demonstrated various approaches for collecting data on occupant attitudes, such as through Likert scales applied to statements about the environment, climate change, energy use, and other topics (Huebner et al., 2016) or through self-identification with political parties or green groups (Wahlström and Hårsman, 2015).

In general, there is limited evidence regarding the effects of occupant attitudes and preferences on electricity consumption patterns. In the few cases where these have been studied, however, findings do not show significant differences between households with different values or attitudes (Brounen et al., 2013; Vringer et al., 2007). Furthermore, these have been shown to explain very little of the variance in electricity consumption in households (Huebner et al., 2016).

Occupant behavioural and activity factors

As was discussed in Section 2.5.1, there are numerous types of occupant behavioural factors, such as those that measure purchasing behaviour, those that indicate habitual energy-saving or curtailment behaviour, and those that represent actual energy-consuming behaviours (e.g., via appliance use, lighting, or cooling and heating activities). Data on energy purchasing and curtailment behaviour are typically collected via household surveys or in-person interviews, and they are included in models as counts of energy efficient appliances purchased or as estimates of the frequency with which an occupant engages in energy curtailment behaviour (e.g., Huebner et al., 2016; Kavousian et al., 2013). Energy-consuming behaviours, specifically appliance use, are typically measured as durations of use, as was discussed previously.

Another method for measuring energy-related behaviours is through the inclusion of time-use data in electricity models (Richardson et al., 2010). Linking time use and electricity use can provide insights on electricity consumption that results both from usage of specific appliances as well as from practices that involve numerous appliances, such as laundry practices or meal preparation (De Lauretis et al., 2017). At first glance, it may seem that collecting both time-use and appliance-use data, especially via monitored appliances, is redundant, but these data do not always align. For instance, Durand-Daubin (2013) found that measured and occupant-reported appliance usage can vary considerably, suggesting systematic bias in the way time-use data are reported. For this reason, empirical validation of time use data in electricity models with measured data is critical (Widén et al., 2009).

Furthermore, time-use data may have important implications for understanding patterns of energy use that cannot be ascertained from monitored or measured data alone. Just because the TV is on, for instance, does not mean that occupants are watching it. This was confirmed empirically in Palmer et al. (2013), who compared monitored appliance data with time-use diaries and found that having the TV on in the background is a common activity.

These are just several of the ways occupant behaviour is accounted for, especially in bottom-up statistical models, but there are numerous approaches for representing occupant energy consumption behaviour in models of electricity demand. A detailed review of these is given in D'Oca et al. (2018).

Evidence for the effects of purchasing or curtailment behaviours on electricity consumption is mixed, with some studies showing these as important for explaining usage patterns (e.g., Brohus et al., 2010) while others show inconclusive findings as to their effect (e.g., Kavousian et al., 2013). Existing research has also found that energy conservation behaviours show much less explanatory power than other categories of factors (Huebner et al., 2016).

Time-use activities have rarely been examined for their explanatory or predictive power in models of residential electricity use. This is primarily a result of the lack of studies that collect both time-use and electricity data in parallel from households, and this is a gap the research presented in this thesis aims to fill.

2.6 Review of key issues and challenges for building energy efficiency and flexibility policy

This section reviews the main issues and challenges relevant to energy efficiency and flexibility policy in order to frame the thesis' third research question on the implications of better understanding the temporal aspects of residential electricity consumption on the design of energy efficiency and flexibility interventions. As was the case with the previous sections, the empirical papers presented in the subsequent chapters provide additional context for the policy implications specific to each paper's findings. In this section, the focus is reviewing the benefits, challenges, and key debates surrounding both building energy efficiency and flexibility as well as discussing existing literature on intervention efforts to change residential electricity demand.

2.6.1 Energy efficiency and demand reduction

Chapter 1 gave a brief introduction to energy efficiency and its importance in achieving ambitious climate change mitigation targets, and Section 2.2 of this chapter summarised some of the recent research exploring the role of energy efficiency in building sector decarbonisation. This section therefore provides a critical review of energy efficiency and demand reduction, discussing the main issues and challenges associated with these concepts.

There is widespread agreement that increasing energy efficiency and reducing energy demand are among the most cost-effective, rapid, and safest options to mitigate climate change (Sorrell, 2015). These mitigation options also offer 'multiple benefits' in the form of improvements to health and well-being, poverty alleviation, economic and social development, and energy security (IEA, 2014). The relationship between increases in energy efficiency and reductions in energy demand is not always straightforward, however, and energy efficiency does not necessarily lead to reduced demand (Sorrell, 2015). Varying approaches to measuring each of these can yield different interpretations of their efficacy. For instance, while energy efficiency is defined as the ratio of useful output given a specific energy input, these inputs and outputs may be measured in purely energy terms (e.g., heat content), in physical terms (e.g., cooking temperature delivered per unit of input), or economic terms (e.g., GDP or economic value per unit of input), and even these terms may not be appropriate for understanding the most relevant output in the energy system—energy service, such as occupant comfort or accessibility, which is more subjective and difficult to measure than the other three (Patterson, 1996; Sorrell, 2015).

Even given these definitional and measurement issues, the broad consensus on energy efficiency has led to shifts among national and international institutions to prioritise policies that deliver reductions in energy demand. For instance, the IEA produces an annual tracker of global trends in energy efficiency, and its 2018 report highlights both that energy efficiency can substantially reduce emissions and energy system costs but also that its full potential is not being realised (IEA, 2018). The report's Efficient World Scenario demonstrates the environmental and economic impacts that would accrue if all countries realised available cost-effective energy efficiency potential. In the buildings sector, the scenario suggests that the average building in 2040 could be nearly 40% more efficient than current buildings in terms of energy use intensity, which would mean global building energy demand could remain flat through 2040, even as total building floor area increases by 60%. Across all sectors, adopting cost-effective energy efficiency measures globally would deliver reductions in energy-related CO_2 emissions equivalent to 40% of the abatement required to meet Paris Agreement targets (IEA, 2018). It is for these reasons that the IEA started producing its annual Energy Efficiency Market Report in 2013—joining existing market reports for oil, coal, gas, and renewable energy—and in doing so acknowledged energy efficiency as a 'first fuel' that governments could invest in at lower costs than other energy sources.

Despite the substantial evidence base that suggests energy efficiency investments are cost effective even without a carbon price (McKinsey & Company, 2009), it is also evident that not all cost-effective opportunities in energy efficiency are taken up (Gerarden et al., 2017). This so-called 'energy efficiency gap' is a well-known challenge. Explanations for its lasting presence are wide-ranging, but in general, they highlight the existence of numerous non-price barriers to adoption (Sorrell et al., 2004). Reducing these costs has been the focus of numerous studies, but given the limitations of traditional economic frameworks in addressing this gap, new approaches from behavioural economics and social psychology are emerging to understand why cost-effective opportunities are not taken up by consumers and how interventions can be designed with this knowledge in mind to increase uptake (Gillingham and Palmer, 2014). This literature is reviewed in the following sections. Still, even as new perspectives engage with the challenge of the energy efficiency gap, it is clear that a multitude of different policy approaches will likely be necessary for energy efficiency to achieve its full potential, especially for more costly and complex solutions beyond just the incremental 'low-hanging fruit' (Rosenow et al., 2017).

In addition to the energy efficiency gap, another key issue related to energy efficiency is the 'rebound effect'. This concept was first noted by the economist William Stanley Jevons, after whom the more general construction of the rebound effect—the Jevons paradox—is named (Jevons, 1866). The rebound effect acknowledges that some of the savings attributable to energy efficiency technology improvements will be lost as a result of increased demand for energy services in response to reduced prices (Herring, 2006). This is the direct rebound effect, whereas the indirect rebound effect refers to the fact that savings due to efficiency can be spent by consumers on other goods and services, thus increasing indirect energy use and emissions. Evidence for the magnitude of the rebound effect is mixed and inconclusive, though a recent meta-review suggests that in developed countries, the direct rebound effect shown in most studies ranges from 5–25% (Gillingham et al., 2016). In other words, if an energy efficiency measure saves 10 kilowatt-hours (kWh) per year, a rebound effect range of 5–25% would reduce these savings by 0.5–2.5 kWh due to consumer responses. Indirect rebound effects are more difficult to measure, as are effects on consumer welfare, macroeconomic price, and macroeconomic growth, though these are nevertheless important components of the total rebound effect.

Moving from the practical to the more theoretical, another critique that has gained prominence in recent years is directed at the concept of energy efficiency itself, which depends on "carving definitions of both energy and service out of the complex interpenetration of everyday technologies and practices" (Shove, 2017, p. 2). This critique takes aim at the abstraction of energy away from important discourses about what energy is used for, which Shove (2017, p. 2) argues is a 'performative' way of thinking about energy efficiency—tied up in discussions focusing mainly on technological improvements—that fails to engage with broader social developments that are responsible for increases in demand. In this way, energy efficiency as a concept is criticised for being so detached from the social aspects of consumption, such as dimensions of practice and notions of comfort, that it fails to meaningfully limit long-term increases in energy demand. In the author's own words, this argument critiques efficiency policies for being "*too* effective, not in reducing demand but in reproducing and stabilising essentially unsustainable concepts of service" (Shove, 2017, p. 7).

Responses to this criticism abound. Fawcett and Rosenow (2017) summarise some of these responses in a commentary on the role of energy efficiency in the low-carbon transition. First, they note that there is strong quantitative evidence that efficiency improvements deliver absolute rather than just relative reductions in demand, citing reports that show efficiency measures were the most important drivers of reduction in U.K. residential energy demand from 2004–2015 (CBER, 2011; CCC, 2017). Second, the authors highlight studies that show maintaining current service levels is not incompatible with radical reductions in emissions (Boardman, 2012; CCC, 2015; National Grid, 2017) and note that evidence to the contrary is lacking in Shove (2017).

This debate is very much an active one in the energy research community, and it will continue to develop as the evidence base on energy efficiency and demand reduction grows. Returning briefly to Section 2.3 and theoretical frameworks for conceptualising energy consumption, it has been argued that too much of the focus in energy efficiency and demand reduction has been on individual choices and technologies rather than on the socio-technical systems that structure and constrain those choices (Geels, 2004). Geels et al. (2017, p. 1242) define socio-technical systems as "the interlinked mix of technologies, infrastructures, organizations, markets, regulations, and user practices that together deliver societal functions such as personal mobility". The authors advocate for a socio-technical approach to understand how these systems influence, for instance, demand for energy services. Such an approach envisions systemic changes that go beyond the adoption of new technologies and instead might include new infrastructure, new markets, new social preferences, and entirely new consumer practices. Understanding how public policy can support and accelerate these transitions is thus a worthwhile endeavour, especially given the need to move beyond incremental progress and achieve more radical, systemic changes to match the scale of the decarbonisation challenge.

2.6.2 Energy flexibility and demand-side response

As was discussed briefly in Chapter 1, energy flexibility and DSR are emerging as complementary solutions alongside energy efficiency, bringing numerous benefits for electricity system costs, reliability, and renewable energy integration (Denholm and Hand, 2011; Junker et al., 2018). As such, these concepts have been studied extensively in recent years. While numerous approaches exist for delivering more flexible demand, such as energy storage, interconnectors, and DSR, this review focuses specifically on building energy flexibility, including both occupant- and technology-led DSR options. This section reviews these areas in the literature, focusing on the benefits, challenges, and key issues related to demand flexibility in buildings.

One of the main benefits of demand flexibility is that it can be used as a resource to provide essential electricity system services at lower costs than alternatives (Aunedi et al., 2013; Bradley et al., 2013; Strbac, 2008). Implementation of demand flexibility can provide services to the grid in the form of generation services, ancillary services, and delivery services (Neukomm et al., 2019). Historically, these essential services have been provided by supply-side entities, such as large-scale generation technologies (Strbac, 2008). Increasingly, however, there is recognition that demand flexibility in buildings can provide these services with much lower capital and investment costs as well as lower environmental and technical risks (National Infrastructure Commission, 2016). Further explanation of each of these grid services and their associated costs is provided below.

Generation services include those for providing energy, which entails power plant fuel, operation, maintenance, and startup/shutdown costs, as well as those for providing system capacity, which entails capital costs for new generating facilities and associated operation and maintenance costs. The costs for building new generation capacity are especially large, indicating substantial cost savings potential for deferring the need to build new power stations (Strbac, 2008). While energy efficiency is likely a more suitable solution for reducing generation operating costs by lowering overall demand, DSR has greater potential to defer generation capacity investments by reducing peak demand through short-term load reductions (Neukomm et al., 2019). Both solutions are useful for lowering generation operating and capacity costs, however, and their cost effectiveness has led to uptake of these resources in forward capacity markets in the U.S. and U.K. in recent years (Liu, 2017).

Ancillary services to the grid include contingency reserves, frequency regulation, and ramping, all of which entail costs related to power plant fuel and operations. These costs are less than generation operating and capacity costs but can still be substantial, especially in power systems that have high penetrations of variable renewable generation (Aunedi et al., 2013). Buildings can provide these services by adjusting demand over short periods of time to provide ancillary services using technologies such as variable frequency drives in heating, ventilation, and air conditioning (HVAC) systems (Hao et al., 2014).

Delivery services are those which entail capital costs for transmission and distribution equipment upgrades as well as voltage control equipment. Capital costs for transmission and distribution upgrades can be significant (though typically not as large as those for new generation capacity) and thus represent a moderate market potential for cost savings. Again, buildings can provide these services in a number of ways (Bradley et al., 2013). In the case of voltage support, however, flexibility from buildings may have difficulty competing with more cost-effective supply-side resources (Neukomm et al., 2019).

In addition to benefits to the grid, DSR can provide benefits to consumers. Perhaps the most obvious of these are the economic benefits accrued via lower utility bills (Aghaei and Alizadeh, 2013; Albadi and El-Saadany, 2008; Strbac, 2008). There is mixed evidence, however, as to the size of these benefits and by how much they improve consumer welfare (see O'Connell et al. (2014) for a review of this evidence). Other benefits mentioned by Neukomm et al. (2019) include improvements to occupant comfort, productivity, and choice, though these are mentioned far less in the literature on the benefits of flexibility.

In the residential sector, there are a range of policy options for DSR, which can be categorised as either price-based or incentive-based (Sharifi et al., 2017). Price-based mechanisms include time-of-use (TOU) pricing, critical-day pricing, critical-peak pricing, and real-time pricing, whereas incentive-based mechanisms include peak-time rebates, demand-side bidding, and 'dynamic demand'. These two categories of options are primarily differentiated by whether they rely on price signals or incentive payments for customers to reduce their electricity use during peak hours or shift it to off-peak hours. The options specific to each range in level of complexity, level of automation or other technical infrastructure required, and level of direct consumer engagement. More details on each option are given in Darby and McKenna (2012) and Strbac (2008).

A further delineation of residential DSR options is whether they are primarily 'appliance-led' or 'activity-led' (Grunewald and Diakonova, 2018). Appliance-led flexibility refers to DSR options that do not necessarily interfere with everyday life, such as automated load reductions in smart appliances, fuel switching in dual-fuel appliances, or features that can shift electricity consumption either later or earlier without requiring user engagement, such as a washing machine with a delayed start feature or a smart thermostat with a pre-cooling function. While appliance-led DSR may appear advantageous in that it requires less action on behalf of the energy user, it has drawbacks related to the potential costs of the technology, loss of user control, and reductions in users' awareness of their energy consumption practices, which may lead to increased usage (Darby and McKenna, 2012). Activity-led flexibility, on the other hand, refers to changes or shifts in consumption patterns that are driven by occupant activities. Such responses might include changing the timing of an activity, substituting an energy-intensive activity with one that does not use as much energy (e.g., reading rather than watching television), fulfilling an energy service with metabolic energy, such as walking or bicycling rather than driving, or changing the practitioner—for instance, using a laundry delivery service rather than washing clothes at home (Grunewald and Diakonova, 2018). These forms of DSR require more behavioural changes on behalf of the energy user and thus incur non-monetary personal and opportunity costs.

While demand flexibility and DSR are not new concepts in the energy demand literature, their implementation has been slowed by several technical, market, and social challenges. These are reviewed in turn. Despite the availability of key DSR-enabling technologies, such as advanced metering infrastructure, wireless communication and control methods, and customer-facing software and interfaces, these information and communications technologies (ICT) are largely absent from electricity systems (Strbac, 2008). Reasons for slow adoption include on-going challenges related to the costs and technical issues involved in integrating these technologies into homes and businesses (O'Connell et al., 2014). Several recent studies have highlighted that investment in these enabling technologies might not be justified given their modest economic benefits for consumers (Gottwalt et al., 2011; Lyon et al., 2012).

In addition to technical challenges, there exist numerous market and regulatory challenges to integrating demand flexibility resources into electricity networks. Establishing a business case for DSR in a market environment where benefits accrue to different participants is one such challenge. Especially in the residential sector, where DSR resources typically must be aggregated to provide a large enough flexibility source for utilities to justify its procurement, there are associated technical and material challenges related to ICT, competition among DSR aggregators, and specification of contracts and remuneration to end consumers (Rajabi et al., 2017). The design of tariff structures for residential customers presents further challenges, as determining appropriate prices or incentives for customer participation in DSR programmes is complex, and even if these are adequately communicated to customers by an energy provider or DSR aggregator, confusion over different options and payment schemes prevents participation (Darby, 2013). Indeed, a recent survey shows that while 14% of all U.S. utilities offer a residential TOU rate, customer enrolment averages around 3% (Hledik et al., 2017). For detailed discussions of additional market and regulatory challenges, interested readers should see O'Connell et al. (2014).

While the technical and market challenges discussed above include social aspects, there are several specific challenges related to the social implications of DSR. First is the issue of end-user behaviour. Whereas conventional systems of electricity provision rely on large generators that typically engage in profit maximising behaviour, DSR instead relies on smaller, distributed end users that do not always exhibit economically rational behaviour (O'Connell et al., 2014). The diversity of end-user behaviour, which has been discussed at length in this and the preceding chapter, makes residential demand a difficult resource to rely on for flexibility purposes. Additional social challenges include the following: residential customers are often reluctant to ask for a time-varying electric tariff if they are used to a standard one, they lack awareness of the rationale and potential benefits of engaging in DSR and, as mentioned previously, there is confusion over the various types of DSR options and how they differ (Kim and Shcherbakova, 2011). Yet another social challenge relates to equity. New tariffs and technologies may prove more or less beneficial for different subsets of the population given the capacity and willingness of different customers to change their behaviours, invest in new technologies, and become more active participants in the electricity system (Darby, 2012). While these challenges are not insurmountable, and while early trials confirm some of these can be addressed by designing informative and engaging customer education programmes and providing useful feedback to customers, key questions remain regarding the social implications of residential DSR.

In part due to the developing understanding that demand flexibility is not purely a technical concept relying on technological fixes but has implications that are socio-technical in nature, recent research has examined residential DSR from a socio-technical perspective to more deeply consider the activities, services, and technologies that are involved (McKenna et al., 2017; Nardelli and Kühnlenz, 2018; Powells et al., 2014). Again, as is the case in the turn towards a socio-technical framework for understanding energy efficiency, there is a need to move beyond simplistic assumptions of how demand-side technologies can provide flexibility and how occupants will respond to (primarily) economic signals to shift demand. Understanding social preferences and consumer practices is essential for achieving the clearly significant benefits and potential of flexibility in the residential sector.

2.6.3 Policy interventions to change demand

As was mentioned previously, one of the obstacles to realising this potential is overcoming the energy efficiency gap, which can also be extended to energy flexibility. Simply put, energy users do not exhibit rational behaviour or optimal decision making in regard to energy use or investments in energy savings. Numerous studies have confirmed this point, noting that energy efficiency products and programmes have seen slower diffusion into the market than should be expected if individuals appropriately valued these products (Gillingham and Palmer, 2014). Related phenomena include the so-called 'knowledge-action gap' and 'value-action gap', which refer to the discrepancy between peoples' knowledge, values, attitudes, and intentions and their actual behaviour (Blake, 1999; Flynn et al., 2010; Huddart Kennedy et al., 2009).

Given these challenges, much research in the behavioural, social, and economic sciences has focused on consumer decision making and behaviour in the context of energy demand. This research has numerous insights for developing energy policies and designing interventions that address this irrational consumer behaviour. Some of this research focuses on identifying those behavioural anomalies and failures that explain why individual behaviour related to energy diverges from that expected from the assumptions of neoclassical economics. These anomalies and failures are categorised by Gillingham and Palmer (2014) as time-inconsistent preferences, such as the preference to invest in energy efficiency products in the future but a reluctance to do so as the future approaches (Tsvetanov and Segerson, 2013); reference-dependent preferences, such as loss aversion (Greene et al., 2009); and non-standard decision making, such as limited attention and use of sub-optimal decision heuristics (Turrentine and Kurani, 2007). Frederiks et al. (2015) review several specific cognitive biases that are especially relevant to energy use: status quo bias, which refers to the tendency among individuals to defer to default settings (Pichert and Katsikopoulos, 2008); risk aversion, which is related to loss aversion but refers to when people exhibit risk averse behaviour in the face of certain gains or uncertain losses but risk seeking behaviour when faced with the opposite; and the sunk cost effect, which describes how people become irrationally persistent with something if they have already invested substantial resources (personal or financial) in it.

These anomalies and failures have long been the focus of behavioural economics, beginning with the pioneering work of renowned psychologists Kahneman and Tversky in the 1970s (Tversky and Kahneman, 1974). This research is relevant to all aspects of consumer behaviour, and it was not originally applied to the case of energy use, but the fact that behavioural anomalies and failures contribute to the energy efficiency gap has been discussed and examined thoroughly in energy research and policy (Gillingham et al., 2009; Shogren and Taylor, 2008; Tietenberg, 2009). A detailed review of these anomalies is beyond the scope of this section; rather, the following review intends to summarise the literature on policies that address these behavioural anomalies and failures in the context of energy demand. The subsequent paragraphs introduce two of these approaches—economic incentives and information provision—and review the intervention options specific to each and the empirical evidence on their effectiveness in relation to energy efficiency and flexibility programmes.

First, economic incentives are commonly used to encourage investments in energy efficiency and reductions in energy demand. As was mentioned in the previous sections, these can be either price-based or incentive-based and include various options for different pricing mechanisms as well as subsidies such as product rebates, tax incentives, and low-cost loans. In this way, these instruments can target overall energy conservation in addition to time-sensitive reductions in energy usage.

Evidence on the effectiveness of dynamic pricing plans as a means to reduce demand is mixed. While Sorrell (2015, p. 79) claims "A mountain of evidence demonstrates that higher energy prices lead to lower energy demand at all levels—as simple theory would predict", Price (2014) is more circumspect in his review of dynamic pricing field experiments in the U.S. His review summarises the findings from pilot experiments conducted at three distinct periods in U.S. history, including very early trials in the 1970s and more recent pilots in the mid-2000s. The findings of this review include the following: reductions in energy use during peak periods among customers on TOU rates are modest because of a highly inelastic short-run demand for electricity (cf. Espey and Espey, 2004); greater responsiveness is observed in homes with some form of AC; critical-peak pricing leads to greater reductions during peak periods (ranging from 13–20%) (Faruqui and George, 2005; Herter et al., 2007), but its effectiveness is dependent on whether the tariff is framed as a rebate or a price, which is consistent with the behavioural anomaly mentioned above related to loss aversion (Wolak, 2011); and finally, these pilots consistently find that responses are stronger when dynamic pricing plans are coupled with enabling technologies, such as real-time displays and smart appliances, that provide feedback and facilitate responses (Jessoe and Rapson, 2012). Interestingly, Faruqui et al. (2010) find through survey responses that in trials that include an in-home display, observed energy savings come from changed behaviours, such as an increased frequency of turning off the lights or large appliances when not in use, rather than purchases of more energy efficient appliances.

Another example of the responsiveness of consumers to electricity pricing comes from studies of California's energy crisis in the early 2000s. Data from this period, which was analysed by Reiss and White (2008), show substantial reductions in average daily household electricity consumption (around 13% for two months) after an unprecedented increase in wholesale electricity prices in California and the western U.S. in June 2000. One of the most interesting findings of this analysis is the speed with which households changed their consumption habits (especially given previously-reported elasticities for electricity demand).

Turning to the evidence on other economic incentives, much of the literature examining the effectiveness of subsidies has focused on utility programmes and, again, shown mixed results. Several studies find that the cost of energy saved is greater than the costs actually claimed by utilities (Arimura et al., 2011; Auffhammer et al., 2008), though the authors in both cases note these differences are not statistically significant.

While most of the earlier research on the design and impact of energy prices was focused on reducing electricity use over the course of weeks and months, more recently there has been a rise in studies that examine the responsiveness of households to time-varying pricing in residential DSR programmes. Srivastava et al. (2018) present a recent review of this literature, comparing 32 studies published since 2006. Their meta-analysis aimed to understand the conditions under which residential DSR programmes are more or less likely to be successful, based on varying definitions of 'success', such as whether they showed statistically significant shifts in load and/or participant savings. The authors find that DSR programmes are more likely to be successful in areas with higher economic growth rates, in urban environments, and in places where renewable energy targets are strong. While they suggest caution in interpreting and generalising from their results given issues related to missing data and homogeneity among programme design, they conclude that the factors determining success of DSR programmes in the studies they reviewed are mostly extrinsic (i.e., having to do with the socio-economic conditions in which the programmes were implemented) rather than intrinsic (i.e., having to do with the specifics of the programme designs themselves). These lessons have implications for DSR policy by suggesting that the success or failure of pricing schemes may have more to do with an enabling political and social environment rather than any particular feature of the programme itself.

As early pricing experiments showed that providing consumers with real-time feedback could increase the effectiveness of pricing policies, more modern trials began to experiment with the use of informational strategies to provide this feedback, which has been found to be complimentary to pecuniary policies to change demand (Price, 2014).

Information strategies are those that provide feedback, tips, and assistance to consumers to identify energy saving opportunities (Abrahamse et al., 2005). Product labels, such as the Energy Star programme run by the U.S. Environmental Protection Agency (EPA), is another example information strategy, as is the use of social norms to provide consumers with feedback on how their electricity use compares to their peers and how they can further reduce their energy use. Information strategies are based on theories that suggest consumers with access to information and feedback on the environmental impacts of their electricity consumption will be more likely to save energy than those without this information (Delmas et al., 2013).

The evidence on the efficacy of these strategies is mixed and suggests different responses depending on the type of information provided and the context in which it is communicated (Delmas and Grant, 2010). Abrahamse et al. (2005) find that information-based interventions on their own have limited effects on consumption. Others, however, find that these programmes can result in much greater energy use reductions (in some cases up to around 30% over 5–8 years) (Ehrhardt-Martinez et al., 2010; Gardner and Stern, 2009). These results are based on different customers, different study contexts, and different types of informational interventions, however, so reliable comparisons of these results are difficult.

In an attempt to quantify the findings on the effectiveness of these strategies, Delmas et al. (2013) perform a comprehensive meta-analysis of results from energy conservation experiments from 1975–2012, which include a total of 156 published trials. They find that households that participated in trials using information-based interventions reduced their electricity consumption by 7.4% on average. The authors also find that providing energy audits, real-time feedback, and individual consultations prove more effective than providing historical or peer comparisons. Further, they find pecuniary incentives, or those that inform individuals of the costs or savings potential associated with their energy use, lead to a relative increase in consumption. This contrasts with the results of Newell and Siikamäki (2013), who find incorporating simple estimates of the economic value of energy savings on the labels of energy efficiency products is the most important factor that guides cost-efficient investments among homeowners. Delmas et al. (2013) conclude their review by claiming that methodological rigour might be lacking in many of the studies they examine, as those reporting the most robust methodologies that control for weather and demographic effects while also using a control group show savings of just 2%.

Regarding the evidence on the effectiveness of social norms, one of the largest experiments that has been conducted was by the company Opower in the U.S. The results showed that including social norms, such as comparisons with neighbours, on household energy bills reduces energy consumption by 2% on average (Allcott and Rogers, 2014). Examining the same programme, Costa and Kahn (2013) find that responsiveness to social norms varies with political ideology, suggesting the need for targeted interventions, while Ayres et al. (2009) find that the use of social norms is especially effective at reducing consumption among high-usage customers.

Despite the somewhat mixed evidence on whether information-based policies work to reduce energy consumption, the research from Delmas et al. (2013) highlights that interventions that target energy savings through changing behaviour, whether by providing feedback or social comparisons, can be effective. This finding suggests that energy efficiency policy should consider these types of interventions alongside economic incentives and technology improvements, especially because they may be relatively low cost in comparison.

Understanding consumer decision making and behaviour in the context of energy use is the subject of decades of research in the behavioural, social, and economic sciences, and this research has yielded important insights for energy policy by identifying specific behavioural anomalies and failures and testing different approaches to address these. While there is, indeed, a mixed evidence base on what approaches work for who and when, there is also agreement that such approaches can change consumer behaviour and, in doing so, improve the cost effectiveness of energy efficiency and DSR programmes. This consensus has important implications for further energy research. First, it suggests that greater awareness of evidence from social and behavioural sciences can inform the proper design of energy efficiency and flexibility interventions, helping to 'bridge the energy efficiency gap' (Gillingham and Palmer, 2014). Second, it further highlights the importance of advancing our understanding of key social and behavioural determinants of electricity consumption, as this knowledge is essential for informing more targeted interventions in the residential sector (Frederiks et al., 2015).

2.7 Chapter summary

This chapter reviewed literature that is relevant to this thesis and its primary research questions. These are the emerging literature on demand-side solutions for climate change mitigation, frameworks for conceptualising residential energy consumption, the methodological literature on electricity modelling approaches, the empirical literature on factors that influence electricity consumption, and policy-relevant literature on energy efficiency and flexibility. As each of the papers includes a thorough literature review, this chapter aimed to give an introduction to the literature that will be reviewed in greater depth in the papers. The chapter has shown that while statistical modelling of electricity consumption and identifying its determinants in the residential sector are both active areas of research, there remain notable gaps, especially regarding the links between occupant activities and electricity consumption. The following chapter presents a methodology for the thesis, showing how the papers aim to fill these gaps by applying new statistical methods to novel sources of data.

2.8 References

Abrahamse, W. and Steg, L. (2009), 'How do socio-demographic and psychological factors relate to households' direct and indirect energy use and savings?', *Journal of Economic Psychology* **30**(5), 711–720.

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.

Abrahamse, W., Steg, L., Vlek, C. and Rothengatter, T. (2007), 'The effect of tailored information, goal setting, and tailored feedback on household energy use, energy-related behaviors, and behavioral antecedents', *Journal of Environmental Psychology* **27**(4), 265–276.

Aghaei, J. and Alizadeh, M.-I. (2013), 'Demand response in smart electricity grids equipped with renewable energy sources: A review', *Renewable and Sustainable Energy Reviews* **18**, 64–72.

Ajzen, I. and Fishbein, M. (1980), *Understanding Attitudes and Predicting Social Behavior*, Prentice-Hall, Englewood Cliffs, NJ.

Albadi, M. H. and El-Saadany, E. F. (2008), 'A summary of demand response in electricity markets', *Electric Power Systems Research* **78**(11), 1989–1996.

Allcott, H. and Rogers, T. (2014), 'The short-run and long-run effects of behavioral interventions: Experimental evidence from energy conservation', *American Economic Review* **104**(10), 3003–3037.

Arimura, T. H., Li, S., Newell, R. G. and Palmer, K. (2011), Cost-effectiveness of electricity energy efficiency programs, Working Paper 17556, National Bureau of Economic Research, Cambridge, MA.

Auffhammer, M., Blumstein, C. and Fowlie, M. (2008), 'Demand-side management and energy efficiency revisited', *The Energy Journal* **29**(3), 91–104.

Aunedi, M., Kountouriotis, P.-A., Calderon, J. E. O., Angeli, D. and Strbac, G. (2013), 'Economic and environmental benefits of dynamic demand in providing frequency regulation', *IEEE Transactions on Smart Grid* **4**(4), 2036–2048.

Ayres, I., Raseman, S. and Shih, A. (2009), Evidence from two large field experiments that peer comparison feedback can reduce residential energy usage, Working Paper 15386, National Bureau of Economic Research, Cambridge, MA.

Baker, K. J. and Rylatt, R. M. (2008), 'Improving the prediction of U.K. domestic energy-demand using annual consumption-data', *Applied Energy* **85**(6), 475–482.

Barr, S., Gilg, A. W. and Ford, N. (2005), 'The household energy gap: Examining the divide between habitual- and purchase-related conservation behaviours', *Energy Policy* **33**(11), 1425–1444.

Bartiaux, F. and Gram-Hanssen, K. (2005), Socio-political factors influencing household electricity consumption: A comparison between Denmark and Belgium, *in* 'Proceedings of the 2005 ECEEE Summer Study on Energy Efficiency in Buildings', European Council for an Energy Efficient Economy, Presqu'île de Giens, France, pp. 1313–1325.

Bartusch, C., Odlare, M., Wallin, F. and Wester, L. (2012), 'Exploring variance in residential electricity consumption: Household features and building properties', *Applied Energy* **92**, 637–643.

Bedir, M., Hasselaar, E. and Itard, L. (2013), 'Determinants of electricity consumption in Dutch dwellings', *Energy and Buildings* **58**, 194–207.

Blake, J. (1999), 'Overcoming the 'value-action gap' in environmental policy: Tensions between national policy and local experience', *Local Environment* **4**(3), 257–278.

Blasch, J., Filippini, M. and Kumar, N. (2019), 'Boundedly rational consumers, energy and investment literacy, and the display of information on household appliances', *Resource and Energy Economics* **56**, 39–58.

Boardman, B. (2012), Achieving zero: Delivering future-friendly buildings, Technical report, Environmental Change Institute, University of Oxford, Oxford, United Kingdom.

Bradley, P., Leach, M. and Torriti, J. (2013), 'A review of the costs and benefits of demand response for electricity in the UK', *Energy Policy* **52**, 312–327.

Brohus, H., Heiselberg, P. K., Simonsen, A. and Sørensen, K. C. (2010), Influence of occupants' behaviour on the energy consumption of domestic buildings, *in* 'Clima 2010: 10th Rehva World Congress: Sustainable Energy Use in Buildings', Clima 2010: 10th Rehva World Congress, Antalya, Turkey.

Brounen, D., Kok, N. and Quigley, J. M. (2012), 'Residential energy use and conservation: Economics and demographics', *European Economic Review* **56**(5), 931–945.

Brounen, D., Kok, N. and Quigley, J. M. (2013), 'Energy literacy, awareness, and conservation behavior of residential households', *Energy Economics* **38**, 42–50.

Brown, M. A. and Li, Y. (2019), 'Carbon pricing and energy efficiency: Pathways to deep decarbonization of the US electric sector', *Energy Efficiency* **12**(2), 463–481.

Brown, M. A. and Macey, S. M. (1983), 'Understanding residential energy conservation through attitudes and beliefs', *Environ Plan A* **15**(3), 405–416.

Carlson, D. R., Scott Matthews, H. and Bergés, M. (2013), 'One size does not fit all: Averaged data on household electricity is inadequate for residential energy policy and decisions', *Energy and Buildings* **64**, 132–144.

Cayla, J.-M., Maizi, N. and Marchand, C. (2011), 'The role of income in energy consumption behaviour: Evidence from French households data', *Energy Policy* **39**(12), 7874–7883.

CBER (2011), British Gas Home Energy Report, Technical report, Centre for Economics and Business Research Limited, London, United Kingdom.

CCC (2015), The fifth carbon budget: The next step towards a low-carbon economy, Technical report, Committee on Climate Change, London, United Kingdom.

CCC (2017), Energy prices and bills - Impacts of meeting carbon budgets, Technical report, Committee on Climate Change, London, United Kingdom.

Chen, J., Wang, X. and Steemers, K. (2013), 'A statistical analysis of a residential energy consumption survey study in Hangzhou, China', *Energy and Buildings* **66**, 193–202.

Costa, D. L. and Kahn, M. E. (2013), 'Energy conservation "nudges" and environmentalist ideology: Evidence from a randomized residential electricity field experiment', *Journal of the European Economic Association* **11**(3), 680–702.

Crawley, D. B., Hand, J. W., Kummert, M. and Griffith, B. T. (2008), 'Contrasting the capabilities of building energy performance simulation programs', *Building and Environment* **43**(4), 661–673.

Creutzig, F., Fernandez, B., Haberl, H., Khosla, R., Mulugetta, Y. and Seto, K. C. (2016), 'Beyond technology: Demand-side solutions for climate change mitigation', *Annual Review of Environment and Resources* **41**(1), 173–198.

Darby, S. J. (2012), 'Metering: EU policy and implications for fuel poor households', *Energy Policy* **49**, 98–106.

Darby, S. J. (2013), 'Load management at home: Advantages and drawbacks of some 'active demand side' options', *Proceedings of the Institution of Mechanical Engineers, Part A: Journal of Power and Energy* **227**(1), 9–17.

Darby, S. J. and McKenna, E. (2012), 'Social implications of residential demand response in cool temperate climates', *Energy Policy* **49**, 759–769.

De Lauretis, S., Ghersi, F. and Cayla, J.-M. (2017), 'Energy consumption and activity patterns: An analysis extended to total time and energy use for French households', *Applied Energy* **206**, 634–648.

Delmas, M. A., Fischlein, M. and Asensio, O. I. (2013), 'Information strategies and energy conservation behavior: A meta-analysis of experimental studies from 1975 to 2012', *Energy Policy* **61**, 729–739.

Delmas, M. and Grant, L. (2010), 'Eco-labeling strategies and price-premium: The wine industry puzzle', *Business & Society* **20**(10), 1–39.

Denholm, P. and Hand, M. (2011), 'Grid flexibility and storage required to achieve very high penetration of variable renewable electricity', *Energy Policy* **39**(3), 1817–1830.

Desmedt, J., Vekemans, G. and Maes, D. (2009), 'Ensuring effectiveness of information to influence household behaviour', *Journal of Cleaner Production* **17**(4), 455–462.

D'Oca, S., Hong, T. and Langevin, J. (2018), 'The human dimensions of energy use in buildings: A review', *Renewable and Sustainable Energy Reviews* **81**, 731–742.

Dolan, P. and Metcalfe, R. D. (2015), 'Neighbors, knowledge, and nuggets: Two natural field experiments on the role of incentives on energy conservation', *Becker Friedman Institute for Research in Economics Working Paper* (2589269).

Durand-Daubin, M. (2013), 'Household activities through various lenses: Crossing surveys, diaries and electricity consumption', p. 13.

Ehrhardt-Martinez, K., Donnelly, K. A. and Laitner, S. (2010), Advanced metering initiatives and residential feedback programs: A meta-review for household electricity-saving opportunities, Technical report, American Council for an Energy Efficient Economy, Washington, D.C.

Espey, J. A. and Espey, M. (2004), 'Turning on the lights: A meta-analysis of residential electricity demand elasticities', *Journal of Agricultural and Applied Economics* **36**(1), 65–81.

Eyre, N. and Killip, G., eds (2019), *Shifting the Focus: Energy Demand in a Net-Zero Carbon UK*, Centre for Research into Energy Demand Solutions (CREDS), Oxford, United Kingdom.

Faruqui, A. and George, S. (2005), 'Quantifying customer response to dynamic pricing', *The Electricity Journal* **18**(4), 53–63.

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

Fawcett, T. and Rosenow, J. (2017), 'Commentary: What is right with energy efficiency? A response to Elizabeth Shove'.

Flynn, R., Bellaby, P. and Ricci, M. (2010), 'The value-action gap in public attitudes towards sustainable energy: The case of hydrogen energy', *Nature, society and environmental crisiseds* pp. 159–180.

Foucquier, A., Robert, S., Suard, F., Stéphan, L. and Jay, A. (2013), 'State of the art in building modelling and energy performances prediction: A review', *Renewable and Sustainable Energy Reviews* **23**, 272–288.

Frederiks, E. R., Stenner, K. and Hobman, E. V. (2015), 'Household energy use: Applying behavioural economics to understand consumer decision-making and behaviour', *Renewable and Sustainable Energy Reviews* **41**, 1385–1394.

Fumo, N. and Rafe Biswas, M. (2015), 'Regression analysis for prediction of residential energy consumption', *Renewable and Sustainable Energy Reviews* **47**, 332–343.

Gardner, G. T. and Stern, P. C. (2009), 'The short list: The most effective actions U.S. households can take to curb climate change', *Environment Magazine* .

Geels, F. W. (2004), 'From sectoral systems of innovation to socio-technical systems: Insights about dynamics and change from sociology and institutional theory', *Research Policy* **33**(6), 897–920.

Geels, F. W., Sovacool, B. K., Schwanen, T. and Sorrell, S. (2017), 'Sociotechnical transitions for deep decarbonisation', *Science* **357**(6357), 1242–1244.

Gerarden, T. D., Newell, R. G. and Stavins, R. N. (2017), 'Assessing the energy-efficiency gap', *Journal of Economic Literature* **55**(4), 1486–1525.

Gillingham, K., Newell, R. G. and Palmer, K. (2009), Energy efficiency economics and policy, Working Paper 15031, National Bureau of Economic Research, Cambridge, MA.

Gillingham, K. and Palmer, K. (2014), 'Bridging the energy efficiency gap: Policy insights from economic theory and empirical evidence', *Rev Environ Econ Policy* **8**(1), 18–38.

Gillingham, K., Rapson, D. and Wagner, G. (2016), 'The rebound effect and energy efficiency policy', *Rev Environ Econ Policy* **10**(1), 68–88.

Gottwalt, S., Ketter, W., Block, C., Collins, J. and Weinhardt, C. (2011), 'Demand side management—A simulation of household behavior under variable prices', *Energy Policy* **39**(12), 8163–8174.

Gram-Hanssen, K. (2009), Introducing and developing practice theory: Towards a better understanding of household energy consumption, *in* 'Proceedings of the Sustaining Everyday Life Conference', Linköping University Electronic Press, Campus Norrköping, Sweden, pp. 45–57.

Gram-Hanssen, K. (2010*a*), 'Residential heat comfort practices: Understanding users', *Building Research & Information* **38**(2), 175–186.

Gram-Hanssen, K. (2010*b*), 'Standby consumption in households analyzed with a practice theory approach', *Journal of Industrial Ecology* **14**(1), 150–165.

Gram-Hanssen, K. (2014), 'New needs for better understanding of household's energy consumption - Behaviour, lifestyle or practices?', *Architectural Engineering and Design Management* **10**(1-2), 91–107.

Gram-Hanssen, K., Kofod, C. and Nœrvig Petersen, K. (2004), Different everyday lives: Different patterns of electricity use, *in* 'Proceedings of the 2004 ACEEE Summer Study on Energy Efficiency in Buildings', American Council for an Energy Efficient Economy, Pacific Grove, California.

Greene, D. L., German, J. and Delucchi, M. A. (2009), Fuel economy: The case for market failure, *in* J. S. Cannon and D. Sperling, eds, 'Reducing Climate Impacts in the Transportation Sector', Springer Netherlands, Dordrecht, Netherlands, pp. 181–205.

Groncow, J. and Warde, A., eds (2001), *Ordinary Consumption*, Routledge, London, United Kingdom.

Grubler, A., Wilson, C., Bento, N., Boza-Kiss, B., Krey, V., McCollum, D. L., Rao, N. D., Riahi, K., Rogelj, J., De Stercke, S., Cullen, J., Frank, S., Fricko, O., Guo, F., Gidden, M., Havlík, P., Huppmann, D., Kiesewetter, G., Rafaj, P., Schoepp, W. and Valin, H. (2018), 'A low energy demand scenario for meeting the 1.5°C target and sustainable development goals without negative emission technologies', *Nature Energy* **3**(6), 515–527.

Grunewald, P. and Diakonova, M. (2018), 'Flexibility, dynamism and diversity in energy supply and demand: A critical review', *Energy Research & Social Science* **38**, 58–66.

Hao, H., Lin, Y., Kowli, A. S., Barooah, P. and Meyn, S. (2014), 'Ancillary service to the grid through control of fans in commercial building HVAC systems', *IEEE Transactions on Smart Grid* **5**(4), 2066–2074.

Hargreaves, T. (2011), 'Practice-ing behaviour change: Applying social practice theory to pro-environmental behaviour change', *Journal of Consumer Culture* **11**(1), 79–99.

Herring, H. (2006), 'Energy efficiency - A critical view', *Energy* **31**(1), 10–20.

Herter, K., McAuliffe, P. and Rosenfeld, A. (2007), 'An exploratory analysis of California residential customer response to critical peak pricing of electricity', *Energy* **32**(1), 25–34.

Higginson, S., Thomson, M. and Bhamra, T. (2014), "'For the times they are a-changin": The impact of shifting energy-use practices in time and space', *Local Environment* **19**(5), 520–538.

Hiteva, R., Ives, M., Weijnen, M. and Nikolic, I. (2018), A complementary understanding of residential energy demand, consumption and services, *in* C. Foulds and R. Robison, eds, 'Advancing Energy Policy: Lessons on the Integration of Social Sciences and Humanities', Springer International Publishing, Cham, pp. 111–127.

Hledik, R., Faruqui, A. and Warner, C. (2017), The national landscape of residential TOU rates: A preliminary summary, Technical report, The Brattle Group, Inc, Boston, MA.

Hsu, D. (2015), 'Identifying key variables and interactions in statistical models of building energy consumption using regularization', *Energy* **83**, 144–155.

Huddart Kennedy, E., Beckley, T. M., McFarlane, B. L. and Nadeau, S. (2009), 'Why we don't "walk the talk": Understanding the environmental values/behaviour gap in Canada', *Human Ecology Review* **16**(2), 151–160.

Huebner, G., Shipworth, D., Hamilton, I., Chalabi, Z. and Oreszczyn, T. (2016), 'Understanding electricity consumption: A comparative contribution of building factors, socio-demographics, appliances, behaviours and attitudes', *Applied Energy* **177**, 692–702.

IEA (2014), Capturing the multiple benefits of energy efficiency, Technical report, International Energy Agency, Paris, France.

IEA (2018), *Energy Efficiency 2018: Analysis and Outlooks to 2040*, OECD/IEA, Paris, France.

Jalas, M. and Juntunen, J. K. (2015), 'Energy intensive lifestyles: Time use, the activity patterns of consumers, and related energy demands in Finland', *Ecological Economics* **113**, 51–59.

Janda, K. B. (2011), 'Buildings don't use energy: People do', *Architectural Science Review* **54**(1), 15–22.

Jessoe, K. and Rapson, D. (2012), Knowledge is (less) power: Experimental evidence from residential energy use, Working Paper 18344, National Bureau of Economic Research, Cambridge, MA.

Jevons, W. S. (1866), *The Coal Question: An Inquiry Concerning the Progress of the Nation, and the Probable Exhaustion of Our Coal-Mines*, 2nd edition edn, Macmillan and Co., London, United Kingdom.

Jin, L., Spurlock, A., Borgeson, S., Fredman, D., Hans, L., Patel, S. and Todd, A. (2016), Load shape clustering using residential smart meter data: A technical memorandum, Technical report, Lawrence Berkeley National Laboratory, Berkeley, CA.

Jones, R. V., Fuertes, A. and Lomas, K. J. (2015), 'The socio-economic, dwelling and appliance related factors affecting electricity consumption in domestic buildings', *Renewable and Sustainable Energy Reviews* **43**, 901–917.

Junker, R. G., Azar, A. G., Lopes, R. A., Lindberg, K. B., Reynders, G., Relan, R. and Madsen, H. (2018), 'Characterizing the energy flexibility of buildings and districts', *Applied Energy* **225**, 175–182.

Kavousian, A., Rajagopal, R. and Fischer, M. (2013), 'Determinants of residential electricity consumption: Using smart meter data to examine the effect of climate, building characteristics, appliance stock, and occupants' behavior', *Energy* **55**, 184–194.

Kim, J.-H. and Shcherbakova, A. (2011), 'Common failures of demand response', *Energy* **36**(2), 873–880.

Knobloch, F., Pollitt, H., Chewpreecha, U., Daioglou, V. and Mercure, J.-F. (2019), 'Simulating the deep decarbonisation of residential heating for limiting global warming to 1.5°C', *Energy Efficiency* **12**(2), 521–550.

Liu, Y. (2017), 'Demand response and energy efficiency in the capacity resource procurement: Case studies of forward capacity markets in ISO New England, PJM and Great Britain', *Energy Policy* **100**, 271–282.

Lyon, T. P., Michelin, M., Jongejan, A. and Leahy, T. (2012), 'Is "smart charging" policy for electric vehicles worthwhile?', *Energy Policy* **41**, 259–268.

McKenna, E., Higginson, S., Grunewald, P. and Darby, S. J. (2017), 'Simulating residential demand response: Improving socio-technical assumptions in activity-based models of energy demand', *Energy Efficiency* .

McKenna, E. and Thomson, M. (2016), 'High-resolution stochastic integrated thermal–electrical domestic demand model', *Applied Energy* **165**, 445–461.

McKinsey & Company (2009), Pathways to a low carbon economy: Version 2 of the global greenhouse gas abatement cost curve, Technical report, McKinsey & Company, New York, NY.

McLoughlin, F., Duffy, A. and Conlon, M. (2012), 'Characterising domestic electricity consumption patterns by dwelling and occupant socio-economic variables: An Irish case study', *Energy and Buildings* **48**, 240–248.

Méjean, A., Guivarch, C., Lefèvre, J. and Hamdi-Cherif, M. (2019), 'The transition in energy demand sectors to limit global warming to 1.5°C', *Energy Efficiency* **12**(2), 441–462.

Mims, N. A., Eckman, T. and Goldman, C. (2017), Time-varying value of electric energy efficiency, Technical Report 1398500, Lawrence Berkeley National Laboratory, Berkeley, CA.

Moberg, K. R., Aall, C., Dorner, F., Reimerson, E., Ceron, J.-P., Sköld, B., Sovacool, B. K. and Piana, V. (2019), 'Mobility, food and housing: Responsibility, individual consumption and demand-side policies in European deep decarbonisation pathways', *Energy Efficiency* **12**(2), 497–519.

Mundaca, L., Ürge-Vorsatz, D. and Wilson, C. (2018), 'Demand-side approaches for limiting global warming to 1.5°C', *Energy Efficiency* .

Nardelli, P. H. and Kühnlenz, F. (2018), 'Why smart appliances may result in a stupid grid: Examining the layers of the sociotechnical systems', *IEEE Systems, Man, and Cybernetics Magazine* **4**(4), 21–27.

National Grid (2017), Future energy scenarios, Technical report, National Grid, Warwick, United Kingdom.

National Infrastructure Commission (2016), Smart Power, Technical report, London, United Kingdom.

Ndiaye, D. and Gabriel, K. (2011), 'Principal component analysis of the electricity consumption in residential dwellings', *Energy and Buildings* **43**(2-3), 446–453.

Nelson, D. J. (2008), Residential baseload energy use: Concept and potential for AMI customers, *in* 'Proceedings of the 2008 ACEEE Summer Study on Energy Efficiency in Buildings', American Council for an Energy Efficient Economy, Pacific Grove, California.

Neukomm, M., Nubbe, V. and Fares, R. (2019), Grid-interactive efficient buildings: Overview, Technical report, U.S. Department of Energy, Washington, D.C.

Newell, R. G. and Siikamäki, J. V. (2013), Nudging energy efficiency behavior: The role of information labels, Working Paper 19224, National Bureau of Economic Research, Cambridge, MA.

Nguyen, A.-T., Reiter, S. and Rigo, P. (2014), 'A review on simulation-based optimization methods applied to building performance analysis', *Applied Energy* **113**, 1043–1058.

O'Connell, N., Pinson, P., Madsen, H. and O'Malley, M. (2014), 'Benefits and challenges of electrical demand response: A critical review', *Renewable and Sustainable Energy Reviews* **39**, 686–699.

Palmer, J., Terry, N. and Kane, T. (2013), Further analysis of the Household Electricity Survey - Early findings: Demand side management, Technical Report 475/09/2012, Cambridge Architectural Research Limited, elementenergy, and Loughborough University, London, United Kingdom.

Palmer, J., Terry, N., Kane, T., Firth, S., Hughes, M., Pope, P., Young, J., Knight, D. and Godoy-Shimizu, D. (2013), Further analysis of the Household Electricity Use Survey - Electrical appliances at home: Tuning in to energy saving, Technical Report 475/09/2012, Cambridge Architectural Research Limited, elementenergy, and Loughborough University, London, United Kingdom.

Patterson, M. G. (1996), 'What is energy efficiency?: Concepts, indicators and methodological issues', *Energy Policy* **24**(5), 377–390.

Pichert, D. and Katsikopoulos, K. V. (2008), 'Green defaults: Information presentation and pro-environmental behaviour', *Journal of Environmental Psychology* **28**(1), 63–73.

Pothitou, M., Varga, L., Kolios, A. J. and Gu, S. (2017), 'Linking energy behaviour, attitude and habits with environmental predisposition and knowledge', *International Journal of Sustainable Energy* **36**(4), 398–414.

Powells, G., Bulkeley, H., Bell, S. and Judson, E. (2014), 'Peak electricity demand and the flexibility of everyday life', *Geoforum* **55**, 43–52.

Price, M. K. (2014), 'Using field experiments to address environmental externalities and resource scarcity: Major lessons learned and new directions for future research', *Oxford Review of Economic Policy* **30**(4), 621–638.

Rajabi, A., Li, L., Zhang, J. and Zhu, J. (2017), Aggregation of small loads for demand response programs – implementation and challenges: A review, *in* '2017 IEEE International Conference on Environment and Electrical Engineering and 2017 IEEE Industrial and Commercial Power Systems Europe', Institute of Electrical and Electronics Engineers, Milan, Italy, pp. 1–6.

Reckwitz, A. (2002), 'Toward a theory of social practices: A development in culturalist theorizing', *European journal of social theory* **5**(2), 243–263.

Reiss, P. C. and White, M. W. (2008), 'What changes energy consumption? Prices and public pressures', *The RAND Journal of Economics* **39**(3), 636–663.

Richardson, I., Thomson, M., Infield, D. and Clifford, C. (2010), 'Domestic electricity use: A high-resolution energy demand model', *Energy and Buildings* **42**(10), 1878–1887.

Røpke, I. (2009), 'Theories of practice — New inspiration for ecological economic studies on consumption', *Ecological Economics* **68**(10), 2490–2497.

Rosenow, J., Kern, F. and Rogge, K. (2017), 'The need for comprehensive and well targeted instrument mixes to stimulate energy transitions: The case of energy efficiency policy', *Energy Research & Social Science* **33**, 95–104.

Sanquist, T. F., Orr, H., Shui, B. and Bittner, A. C. (2012), 'Lifestyle factors in U.S. residential electricity consumption', *Energy Policy* **42**, 354–364.

Santin, O. G., Itard, L. and Visscher, H. (2009), 'The effect of occupancy and building characteristics on energy use for space and water heating in Dutch residential stock', *Energy and Buildings* **41**(11), 1223–1232.

Schatzki, R, T. (1996), *Social Practices: A Wittgensteinian Approach to Human Activity and the Social*, Cambridge University Press, Cambridge, United Kingdom.

Schatzki, R, T. (2002), *Site of the Social. A Philosophical Account of the Constitution of Social Life and Change*, Pennsylvania State University Press.

Sharifi, R., Fathi, S. H. and Vahidinasab, V. (2017), 'A review on demand-side tools in electricity market', *Renewable and Sustainable Energy Reviews* **72**, 565–572.

Shogren, J. F. and Taylor, L. O. (2008), 'On behavioral-environmental economics', *Rev Environ Econ Policy* **2**(1), 26–44.

Shove, E. (2010), 'Beyond the ABC: Climate change policy and theories of social change', *Environ Plan A* **42**(6), 1273–1285.

Shove, E. (2017), 'What is wrong with energy efficiency?', *Building Research & Information* pp. 1–11.

Shove, E. and Walker, G. (2010), 'Governing transitions in the sustainability of everyday life', *Research Policy* **39**(4), 471–476.

Shove, E. and Walker, G. (2014), 'What Is energy for? Social practice and energy demand', *Theory, Culture & Society* **31**(5), 41–58.

Sonnenschein, J., Van Buskirk, R., Richter, J. L. and Dalhammar, C. (2019), 'Minimum energy performance standards for the 1.5°C target: An effective complement to carbon pricing', *Energy Efficiency* **12**(2), 387–402.

Sorrell, S. (2015), 'Reducing energy demand: A review of issues, challenges and approaches', *Renewable and Sustainable Energy Reviews* **47**, 74–82.

Sorrell, S., O'Malley, E., Schleich, J. and Scott, S. (2004), *The Economics of Energy Efficiency: Barriers to Cost-Effective Investment*, Edward Elgar Publishing, Cheltenham, United Kingdom.

Sovacool, B. K. (2014), 'What are we doing here? Analyzing fifteen years of energy scholarship and proposing a social science research agenda', *Energy Research & Social Science* **1**, 1–29.

Srivastava, A., Van Passel, S. and Laes, E. (2018), 'Assessing the success of electricity demand response programs: A meta-analysis', *Energy Research & Social Science* **40**, 110–117.

Steemers, K. and Yun, G. Y. (2009), 'Household energy consumption: A study of the role of occupants', *Building Research & Information* **37**(5-6), 625–637.

Strbac, G. (2008), 'Demand side management: Benefits and challenges', *Energy Policy* **36**(12), 4419–4426.

Swan, L. G. and Ugursal, V. I. (2009), 'Modeling of end-use energy consumption in the residential sector: A review of modeling techniques', *Renewable and Sustainable Energy Reviews* **13**(8), 1819–1835.

Tietenberg, T. (2009), 'Reflections - energy efficiency policy: Pipe dream or pipeline to the future?', *Rev Environ Econ Policy* **3**(2), 304–320.

Torriti, J. (2014), 'A review of time use models of residential electricity demand', *Renewable and Sustainable Energy Reviews* **37**, 265–272.

Tso, G. K. and Yau, K. K. (2007), 'Predicting electricity energy consumption: A comparison of regression analysis, decision tree and neural networks', *Energy* **32**(9), 1761–1768.

Tsvetanov, T. and Segerson, K. (2013), 'Re-evaluating the role of energy efficiency standards: A behavioral economics approach', *Journal of Environmental Economics and Management* **66**(2), 347–363.

Turrentine, T. S. and Kurani, K. S. (2007), 'Car buyers and fuel economy?', *Energy Policy* **35**(2), 1213–1223.

Tversky, A. and Kahneman, D. (1974), 'Judgment under uncertainty: Heuristics and biases', *Science* **185**(4157), 1124–1131.

Vringer, K., Aalbers, T. and Blok, K. (2007), 'Household energy requirement and value patterns', *Energy Policy* **35**(1), 553–566.

Wachsmuth, J. and Duscha, V. (2019), 'Achievability of the Paris targets in the E.U. - The role of demand-side-driven mitigation in different types of scenarios', *Energy Efficiency* **12**(2), 403–421.

Wahlström, M. H. and Hårsman, B. (2015), 'Residential energy consumption and conservation', *Energy and Buildings* **102**, 58–66.

Walker, G. (2014), 'The dynamics of energy demand: Change, rhythm and synchronicity', *Energy Research & Social Science* **1**, 49–55.

Wallis, H., Nachreiner, M. and Matthies, E. (2016), 'Adolescents and electricity consumption: Investigating sociodemographic, economic, and behavioural influences on electricity consumption in households', *Energy Policy* **94**, 224–234.

Widén, J., Lundh, M., Vassileva, I., Dahlquist, E., Ellegård, K. and Wäckelgård, E. (2009), 'Constructing load profiles for household electricity and hot water from time-use data—Modelling approach and validation', *Energy and Buildings* **41**(7), 753–768.

Wilhite, H., Shove, E., Lutzenhiser, L. and Kempton, W. (2000), Twenty years of energy demand management: We know more about individual behavior but how much do we really know about demand, *in* 'Proceedings of the 2000 ACEEE Summer Study on Energy Efficiency in Buildings', American Council for an Energy Efficient Economy, Pacific Grove, California, pp. 435–453.

Wilson, C., Pettifor, H., Cassar, E., Kerr, L. and Wilson, M. (2018), 'The potential contribution of disruptive low-carbon innovations to 1.5°C climate mitigation', *Energy Efficiency* .

Wolak, F. A. (2011), 'Do residential customers respond to hourly prices? Evidence from a dynamic pricing experiment', *American Economic Review* **101**(3), 83–87.

Wyatt, P. (2013), 'A dwelling-level investigation into the physical and socio-economic drivers of domestic energy consumption in England', *Energy Policy* **60**, 540–549.

Chapter 3

Methodology

3.1 Introduction

This chapter has two primary aims. First, it provides an introduction to statistical learning and an overview of the methods used in the thesis. Second, it explains the structure of analysis undertaken in the research and how the individual papers form a coherent and complete whole. Section 3.2 first provides several practical details involved in conducting this three-year research project. Section 3.3 then gives an introduction to statistical learning, including a high-level overview of the different types of methods that are applied in this thesis. This section does not delve deeply into explaining the methods because the individual papers have detailed methods sections. Section 3.4 provides a summary of and justification for the methods used. Finally, Section 3.5 explains the structure of the analysis and briefly describes the four research papers, including their specific methodological, empirical, and policy research questions.

3.2 Practical details

3.2.1 Development of work

The focus and approach of this thesis has changed course several times during the past three years. I was admitted to the DPhil through Oxford's 2+2 MPhil-to-DPhil scheme and, as such, the work I undertook to complete my MPhil thesis formed the basis for the first paper of my DPhil. The published version of this paper, however, ended up being substantially different in both scope and methodology after peer review. It was not until the next year, the second of my DPhil, that I began to settle on the overall topic for the thesis of characterising patterns of electricity demand at different temporal resolutions using statistical learning methods. During my second year, I mostly worked on research related to the Measuring and Evaluating Time- and Energy-use Relationships (METER) project, a 5-year Engineering and Physical Sciences Research Council project led by my supervisor. This research focused on applying the methods I developed in my first paper to these new time- and electricity-use data. This work, which carried on into my third year, led to papers 2 and 3 of this thesis.

3.2.2 Research assistant positions

Between my second and third years, I had the opportunity to spend the summer as a research assistant at the U.S. Department of Energy's Building Technologies Office, where I was employed by Lawrence Berkeley National Laboratory. It was there that I designed and carried out the research presented in paper 4. My agreement with my employer was that I would contribute to the group's data processing, visualisation, and analysis efforts so long as I could pursue my own research project during the months I was there with the aim to publish my findings in a peer-reviewed journal. The analytical work took place during the months I was formally employed there, but the paper was written in the final year of my DPhil. I also worked part-time as a research assistant on the METER project during my third and final year. I was responsible for managing and recruiting study participants, processing research equipment for time-use and energy-use data collection, handling and querying project data, and contributing towards the development of research and analytical methods. This work was in addition to the specific analyses undertaken for papers 2 and 3.

3.2.3 Conference attendance

Two conferences have been instrumental to the development and distribution of this research. The first was the American Council for an Energy Efficient Economy (ACEEE) 2018 Summer Study on Energy Efficiency in Buildings, at which I presented a much-abbreviated version of paper 2 of this thesis. After the conference, I was notified that the paper had been selected as one of several to be invited for publication in a special issue of the journal *Energy Efficiency*, and this publication is the second paper in this thesis. Next, I presented a shortened version of my third paper at the European counterpart to the ACEEE Summer Study.[1] As was the case with paper 2, much further analysis and revision was undertaken before submitting this paper for publication.

3.2.4 Ethical permission

As the first three papers involve research on human subjects, appropriate ethical permission was granted for these by the Central University Research and Ethics Committee of the University of Oxford. The approval process for the first paper was particularly time-consuming because it required a full non-disclosure agreement and accompanying privacy policies to be signed between the University and the City of Palo Alto Utilities for data access and storage. For each of the first three papers, all participant data have been anonymised and are held on secure servers. As explained in the papers, a key output of this research is publicly accessible data for use by the energy research community. While the data for paper 1 have tighter restrictions on use and sharing (available only upon request), the anonymised data for papers 2 and 3 are publicly available in a U.K. Data Service repository. The accompanying data collection and processing tools have also been made available to the research community. Finally, the simulation tool used in paper 4 is open access, and the data applied in the analyses are also available upon request.

3.3 Overview of statistical learning and its application in this thesis

As will be made clear in this section, many of the methods used throughout this thesis are, for the most part, new in terms of their application to household and electricity consumption data. For this reason, and due to peer-review comments received in the early stages of publishing this research, the papers included in this thesis (with the exception of the fourth

[1] European Council for an Energy Efficient Economy (ECEEE) 2019 Summer Study on Energy Efficiency in Buildings.

paper) are somewhat longer than typical journal publications and include comprehensive methods sections that fully explain the statistical learning methods applied in each. This section thus provides only a high-level overview of the field of statistical learning and the sub-fields which are the focus of this research.

An important note to the reader concerns the use of terminology in this thesis. While most will be familiar with the term 'machine learning', some may be less familiar with 'statistical learning' and how it differs. For the most part, the terms are used interchangeably to refer to tools that facilitate learning from data. Both approaches are used by practitioners to understand patterns in data, including the structures and relationships between data inputs and outputs. Some of the differences between the two include the following: 1) the broad disciplines they fall under, where machine learning is more commonly associated with computer science and artificial intelligence while statistical learning is more commonly associated with mathematics; 2) the type of modelling approach they use, where machine learning often refers to 'deep learning' and other non-linear methods while statistical learning more broadly refers to both linear and non-linear methods; 3) the type of data used as inputs, where machine learning, due to the use of non-linear models, usually requires millions of data points while statistical learning approaches can be used with smaller data sets; and, finally, 4) the nomenclature specific to each. This last point may be the most noticeable difference between the two, as the elements of machine learning typically use different naming conventions than those of statistical learning, even though they generally refer to the same thing (e.g., 'weights' in machine learning versus 'parameters' in statistical learning).

It was intentional to use the term 'statistical learning' throughout this thesis for the following reasons: first, while many different methods are applied in the papers, these are mostly linear methods, albeit with more advanced techniques applied for sparsity (see Section 3.3.2). Second, the data sets used in the papers have observations in the hundreds to thousands rather than millions. Third, as a result of using linear methods with regularisation, one of the key aims of the modelling exercises undertaken is interpretability (see Section 3.3.1). For these reasons, the general nomenclature used throughout the thesis is that used in statistical disciplines.[2]

3.3.1 Two applications for statistical learning: prediction and inference

As statistical learning refers to tools that are used to better understand, characterise, and learn from data, all statistical learning methods are designed to map input variables, which are variously called predictors, independent variables, or features, to an output variable, called the response or dependent variable, via some unknown function f based on a set of observed data points. The aim of statistical learning approaches is to estimate the unknown function f based on the observed data.

Estimating f from the observed data is done for two main reasons: prediction and inference. While these are not mutually exclusive and can be pursued in tandem, the choice of specific method or algorithm is usually guided by whether the estimation task is primarily for prediction or inference.

Prediction refers to situations where the primary concern is how well our estimate for f, which we represent as $\hat{f}$, can predict new values of a response variable based on the input variables. These situations generally arise when the set of inputs are easily accessible but the output is not. Consider the following simple example: if a building manager wants to accurately forecast day-ahead building electricity demand in the afternoon hours for the purpose of participating in a DSR programme, she might do so by learning a function $\hat{f}$ to predict building electricity demand based on various inputs such as hourly temperature and humidity. In this context, we assume temperature and humidity data can easily be measured and are available as inputs to our statistical learning algorithm. Because the primary interest is predicting building electricity demand, which may then subsequently be decreased or increased via some automated building control system, the chief concern is how well the algorithm performs in terms of its predictive accuracy. The building engineer likely

[2]See James et al. (2013) for a recommended reference on statistical learning that includes its terminology and methods.

cares less about the exact form of $\hat{f}$ or, in other words, *how* it makes predictions based on inputs; instead, she simply wants to know whether the function can accurately inform an automated building management system whether or not to schedule a turn-down for the cooling system the following afternoon. In these contexts, knowing what is happening 'under the hood' is not crucial.

The second reason to estimate f is not necessarily to make accurate predictions but rather to understand how the output changes as a function of the inputs. In this situation, knowing the exact form of $\hat{f}$ is essential. Returning to our previous example, if the main interest is understanding how the temperature and humidity inputs are associated with building electricity demand, for instance for the purpose of identifying which of these is more important for understanding variations in electricity demand, an inferential modelling approach is necessary. A policymaker might want to know how expected increases in regional temperatures and humidity levels will affect building electricity demand. It may be the case that temperature is positively associated with electricity demand but another input variable, the efficiency rating of the building, is negatively associated with demand. Understanding *how* these various inputs relate to the output variable is important for subsequently attempting to *change* it.

The papers in this thesis are concerned with both predictive accuracy and with the relationships between inputs and outputs, but the main motivation behind the methods applied in the thesis is to assess the relationship between many different types of structural and occupant-related factors and electricity consumption at different temporal resolutions. For this reason, it should be made clear that predictive accuracy is not considered the primary objective in the modelling exercises undertaken. That is not to say it is not considered, however, and in each of the papers that apply statistical learning techniques, the results always include a discussion of model predictive accuracy or fit.

Because the main motivation of this research is to understand the relationships between different input variables and electricity consumption, as well as how these inputs compare to one another, linear models are used throughout. In comparison to non-linear methods, linear methods are less flexible and more restrictive, meaning the functional relationship between inputs and outputs is constrained to be linear in shape. Non-linear models, on the other hand, are more flexible and less restrictive, allowing $\hat{f}$ to take on more complex, non-linear shapes. One might wonder why a modeller would use a less flexible approach when more flexible ones exist. There are several reasons, but the most relevant here is because the less flexible linear models are more interpretable than the more flexible non-linear models. For very flexible approaches, the estimates of f can be so complicated that there is no practical significance to the relationships between individual predictors and the response. This form of the interpretability challenge, where it is challenging or even impossible to estimate the mathematical relationships between inputs and outputs, is a key analytical challenge in statistical learning when measuring these relationships is important (Shmueli, 2010). Furthermore, the highly flexible models also run the risk of a statistical challenge called 'overfitting', which occurs when a model is fit so perfectly to observed data that it does not generalise to new data and thus suffers in terms of predictive accuracy as a result. A more detailed discussion of the issue of overfitting in the context of energy modelling is given in the first paper.

Returning to the question of whether the statistical learning objective is prediction or inference, it is advantageous to use less flexible methods when the aim is inference, in which cases it is important to know the relationships between individual predictors and the response variable. When there are many potential predictive variables that might be associated with the response, however, even simple linear regression models can quickly become uninterpretable. This is the second challenge of interpretability that was referred to in Chapter 1. This issue is prevalent when dealing with large, multivariate data sets that are common in energy research. As will be discussed in the following section, there are a class of methods that can be applied to linear regression that perform variable selection to yield a sparser model, which is slightly less flexible than the OLS model but more interpretable. These are called regularisation methods, and they enable statistical learning with sparsity.

3.3.2 Statistical learning with sparsity: model selection and regularisation

The first paper included in this thesis provides a detailed overview of regularisation and its use to achieve sparse statistical models that benefit from more interpretability without sacrificing predictive accuracy. While this overview should be sufficient to guide the reader through the subsequent papers, a brief, general overview of statistical learning with sparsity is provided here.

In the context of linear regression, regularisation, which is also sometimes called a shrinkage method or penalised regression, is a technique that is used to constrain or 'regularise' the coefficients of a linear model's parameters by shrinking these toward zero. This is, in actuality, a biased estimate for the function $\hat{f}$ rather than the true least squares estimate, but when regularisation is applied, the aim is to trade off a small increase in bias for a much larger reduction in variance.[3] As a result of this reduction in variance, the model is less likely to overfit the data and thus will make better predictions on new, unseen data.

But beyond this benefit, there is another key element to some regularisation methods, such as lasso and elastic net regression, and this is their ability to deliver sparse solutions. A sparse statistical model has only a small number of non-zero parameters due to the regularisation methods shrinking many of the other parameters' coefficients to zero (Hastie et al., 2015). A sparse model is advantageous because it is easier to interpret than a dense model. Especially in situations where there is a large set of possible predictors that potentially relate to the response variable and should be included in the model, the use of sparse statistical learning methods simplifies these modelling exercises and enables researchers and practitioners to extract the most useful information from these data.

Another advantage to these methods is that they can be used to estimate f in situations where other conventional methods fail. One example of this is when the number of predictors p is larger than the number of observations n. In these cases, OLS will fail to yield a unique

[3] A detailed explanation of the concepts of bias and variance is given in the first paper.

solution and thus cannot be used. Several regularisation methods do, however, work in these situations, which opens up new avenues in modelling contexts that suffer from these so-called ‘high-dimensional’ data.

There are clear benefits for using regularisation methods in statistical models of electricity demand. Variable selection, especially, is key when the aim of the modelling exercise is to better understand the factors that influence consumption. And in the incredibly rich built environment, where an almost limitless number of factors exist that indicate how buildings are constructed, lived in, and, as a result, consume electricity, sparser models are essential for capturing key variables that have important ramifications for understanding electricity use.

3.3.3 Supervised and unsupervised learning

There are two main categories of statistical learning approaches: supervised and unsupervised. In the supervised learning context, the data are labelled such that each observation has measurements for both the predictor variables and the response variable. As explained previously, the statistical learning task is to estimate a function $\hat{f}$ that relates the response to the predictors, and the aim of this task can either be to accurately predict the response given new observations or to mathematically model the relationship between the two. Most well-known statistical learning methods exist for the supervised learning setting. Figure 3.1 presents a graphical overview of the types of supervised and unsupervised learning algorithms described below.

Supervised learning problems are classified as either regression (Figure 3.1a) or classification (Figure 3.1b) problems. In the regression setting, the response variable is characterised as a quantitative variable in that it takes on numerical values. A relevant example is electricity consumption, which takes on a numerical value, usually measured in Watt-hours (Wh). Whether the electricity usage is recorded over seconds, hours, or years, it is still a quantitative variable and so is usually estimated using regression. Classification problems, on the other hand, have response variables that are qualitative, or categorised. These take

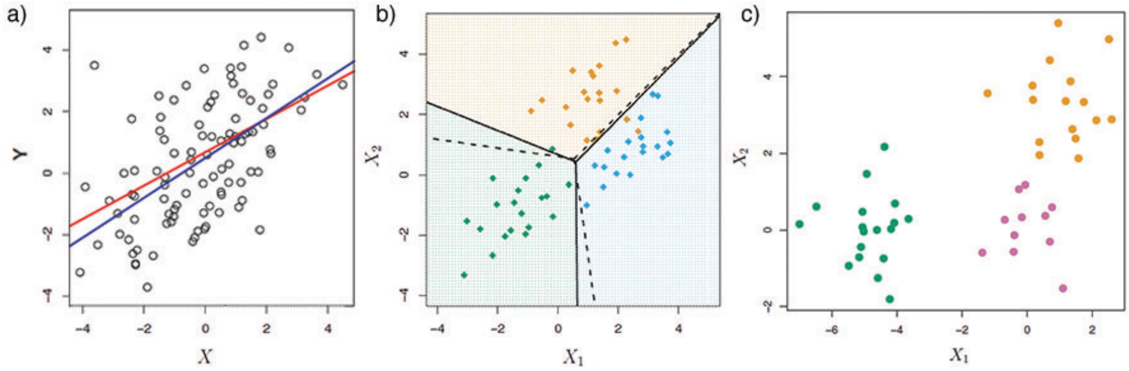

Figure 3.1: Graphical representation of different types of statistical learning problems. (a), Supervised learning (regression). (b), Supervised learning (classification). (c), Unsupervised learning (clustering). Adapted from James et al. (2013).

on values that represent different classes or categories. A quantitative variable measuring electricity consumption could be converted to a qualitative variable by assigning any value below a certain threshold as 'low' and any value above that threshold as 'high'. These two classes form a binary response variable. As illustrated by this example, the distinction between regression and classification problems is not always distinct (for instance, one of the most common classification methods, logistic regression, is used with a binary response variable but can also be thought of as a regression method because it estimates class probabilities).

Unsupervised learning refers to contexts in which every observation has measurements for its associated predictor variables, but there is no label or measurement for the associated response variable. In this situation, there is neither a quantitative response variable to predict using regression nor a qualitative class to predict using classification. The term unsupervised refers to the fact that there is no response variable to 'supervise' the analysis. In cases where response variables are not labelled or measured for a given data set, there is still useful knowledge we can learn from the data. For instance, we can understand the relationships not between the predictors and the response but between the predictors themselves. Say, for example, that a utility programme designer wants to group residential customers by various indicators such as their monthly utility bill, their zip code, and other related variables in order to group customers into distinct segments that can be used to improve the effectiveness of a utility programme. In this case, the customers can be grouped together using an unsupervised learning algorithm such as cluster analysis, which uses mathematical relationships between customers based on their characteristics to distinguish which are more similar and should thus be grouped together.

This thesis applies both types of statistical learning approaches as well as both types of supervised learning methods in the empirical papers. As will be discussed in Chapter 8, while in most cases the application of these different types of approaches is clear cut and guided by past research, in some cases the statistical learning problem can be set up as either supervised or unsupervised and as either regression or classification. In instances where there was the option to choose which approach to use, the research was guided by a general interest in applying different types of approaches and algorithms to showcase the suitability of many statistical learning methods for characterising electricity demand.

3.3.4 Resampling methods

A final category of methods that are integral to the modelling efforts undertaken in this thesis are resampling methods. These methods involve repeatedly resampling a data set and refitting a model in order to obtain additional information, such as the variability of the model fit, which could not otherwise be ascertained if resampling methods were not used. Because they essentially involve repeating any given analysis multiple (and often many) times using different resamples of the same data, these methods can be computationally intensive. But most computational software available today can relatively quickly resample data sets and refit a model or parameter of interest hundreds of times. There are two resampling methods that are applied in this thesis, cross-validation and bootstrap resampling, and both are briefly reviewed in turn.

Cross-validation refers to any number of different approaches to hold out a subset of the data a model is trained on, which can then be used to estimate the error rate the model is likely to have given new data. In the simplest case, the 'model validation set' approach involves randomly splitting the data into a training set and a validation set, and after the model is fit to the training set, the validation set is used to estimate the test set error. When only one randomly sampled portion of the original training data is used to validate a model, however, the estimate for its error can be highly dependent on the specific observations that are included in the sample. For this reason, cross-validation is introduced to repeatedly draw these resamples to get a less variable estimate for the model's error.

There are several different types of cross-validation, including leave-one-out cross-validation (LOOCV), leave-p-out cross-validation (LpO-CV), and k-fold cross-validation. In the papers in this thesis, k-fold cross-validation is used. Briefly, this method involves randomly splitting the training data into k-folds, training the model on $k-1$ folds and using the held out fold to validate the model, and then repeating the process until each of the k folds has served as the validation set. The resulting k estimates of model error are averaged to arrive at an overall cross-validation error. Typically, a value of $k = 5$ or $k = 10$ is used, as these have empirically been shown to yield low test set error rates in most applications (James et al., 2013). Figure 3.2 presents a graphical overview of k-fold cross-validation.

In the case of models that have their own adjustable parameters, which are called hyperparameters in statistical and machine learning, cross-validation is a useful approach for tuning these for the purpose of finding or selecting the model that performs the best across a range of different values for the hyperparameters. Regularisation methods require hyperparameter tuning, and cross-validation is an integral part of this procedure. The first three papers each apply cross-validation for this purpose and so provide a more detailed explanation of the tuning process.

Bootstrap resampling is another powerful statistical tool for quantifying the uncertainty of a given model or method. The bootstrap process enables modellers to computationally emulate the process of obtaining new data sets in order to estimate the variability of a quantity of interest. It does this by resampling observations from the original data set in order to produce a bootstrap data set. This resampling is done with replacement, meaning

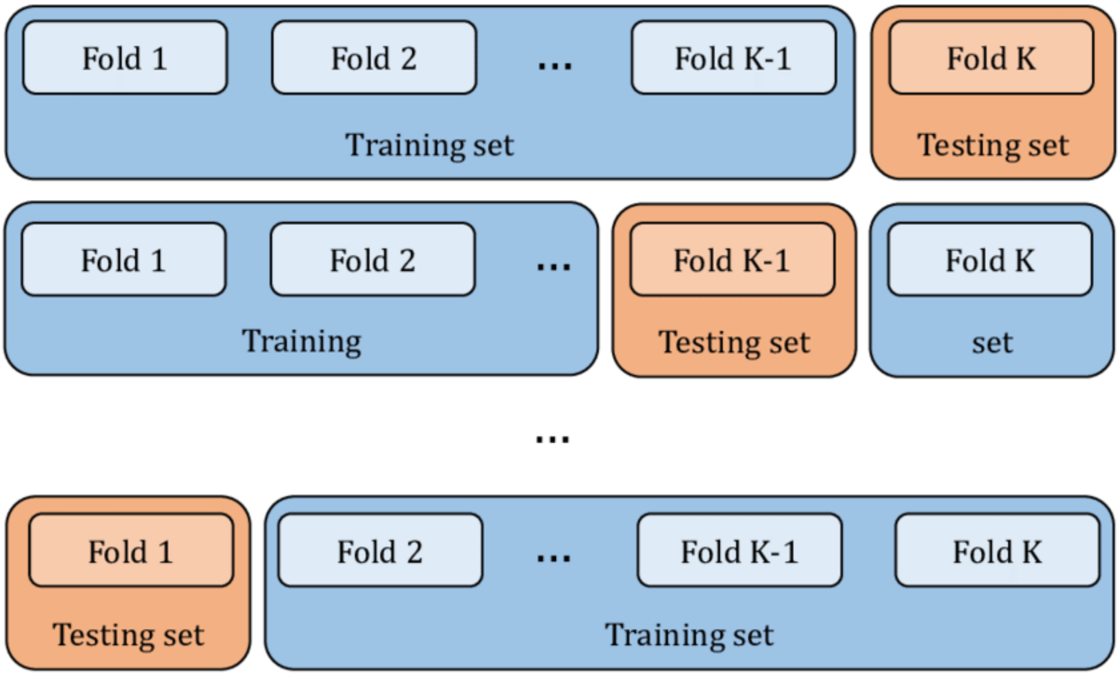

Figure 3.2: Overview of the k-fold cross-validation procedure for obtaining an overall estimate of training error. Adapted from Viegas et al. (2016).

any observation can be sampled more than once in each bootstrap resample. This procedure is repeated any number of times (often thousands) to create B bootstrap data sets and corresponding estimates, for instance for the regression coefficients of a linear model. If B is large enough, bootstrapping can give an accurate estimation of variability. This is especially useful in cases where an estimate of variability is difficult to obtain from a particular statistical learning method.

Bootstrap resampling is integral to one of the statistical learning methods that is used in paper 3. Its application in this context highlights its usefulness for ensuring the results of a particular model are not simply due to the inclusion of specific observations in the training data by random chance and instead are robust to numerous resamples of the data.

3.4 Justification of methods used

The methods employed throughout this thesis go beyond conventional regression approaches for modelling electricity demand by applying a variety of regularisation techniques to achieve sparse models that show improved accuracy in prediction tasks while also performing variable selection. This second feature means these methods can simplify final regression models by selecting only variables that have the strongest associations to the dependent variable. In yielding sparse models, these techniques prove useful for considering a wide variety of structural and occupant variables and identifying those that are key to explaining residential electricity use at different temporal resolutions.

A detailed discussion of the advantages of regularisation methods in the context of energy research is provided in the first paper of this thesis and returned to in Chapter 8, but given that the thesis' methodological research question asks how such methods "improve efforts to characterise residential electricity consumption patterns", some of this discussion is also included here. This section justifies the use of regularisation methods in this thesis by providing an overview of how they compare to other types of statistical learning approaches, including conventional linear and multiple regression, subset selection methods, and non-linear or 'deep learning' methods. Table 3.1 presents a summary of these different approaches, their common methods, and the positive and negative attributes of each approach. The subsequent paragraphs explain how and under what circumstances regularisation methods are expected to 'improve efforts' to characterise residential electricity consumption.

Linear regression approaches have advantages in terms of their straightforward approach to fitting, prediction, and inference (James et al., 2013). The coefficients, standard errors, confidence intervals, and p values obtained via OLS can generally be trusted, so long as the assumptions on which these statistical tests rely are valid (Osborne and Waters, 2002). As explained in Section 3.3.1, OLS performs well for both prediction and inference when the true relationship between predictors and the response is approximately linear. These

	Linear regression	**Subset selection**	**Regularisation**	**Non-linear methods**
Common methods	• Simple linear regression • Multiple linear regression	• Best subset selection • Stepwise selection	• Ridge regression • Lasso regression • Elastic net regression	• Tree-based methods • Support vector machines • Artificial neural networks
Positive attributes	• Easy to understand and implement • Predict well if true relationship is approximately linear • Inference is typically straightforward (coefficients are inherently interpretable)	• Conceptually simple and straightforward • Perform variable selection • Can prevent overfitting by indirectly or directly estimating test error via cross validation (C_p, AIC, BIC)	• Sparse methods perform variable selection • Can fit high-dimensional data • Can prevent overfitting with large enough penalty • Convex property eases computational complexity • Can handle multicollinearity	• Highly flexible • Predict better than other methods when true relationship is non-linear • Work well in high-dimensional situations • Can handle many different data types
Negative attributes	• Infinitely many solutions in high dimensions • Unlikely to yield coefficient estimates that are exactly zero (does not perform variable selection) • Severe issues handling multicollinearity • Prediction accuracy suffers if $n \approx p$	• Best subset selection is non-convex (can be computationally intractable) • Biased coefficient estimates • R^2 values biased upwards • Artificially low standard errors and narrow confidence intervals • Issues handling multicollinearity	• Ridge does not perform variable selection • Lasso can select at most n predictors in high-dimensional problems • Lasso arbitrarily selects multicollinear predictors • Inference is not straightforward	• Can be conceptually challenging • Can easily overfit the training data • Not interpretable (coefficients lack physical significance) • Can be computationally difficult • Require large training data sets

Table 3.1: Overview of different statistical learning approaches, specific methods, and the positive and negative attributes of each.

both suffer, however, as the number of predictors approaches the number of observations and when multicollinearity effects are present (Hsu, 2015). In high-dimensional problems, when the number of predictors p is greater than the number of observations n, OLS cannot be used because there are infinitely many solutions to the least squares fit (James et al., 2013). Furthermore, a key drawback of OLS is that it cannot perform variable selection.

At the other end of the spectrum, non-linear methods work well in high-dimensional situations, they can handle many different data types, and they often have high predictive accuracy for modelling complex phenomena given their flexibility (James et al., 2013). But this flexibility makes them prone to overfitting and far less interpretable than linear regression methods.

Given this thesis' dual objectives of prediction accuracy and model interpretability in analyses of residential electricity consumption, both conventional linear regression techniques and complex, non-linear methods are not as well suited to these analyses because of the two interpretability challenges mentioned previously. The inability of OLS to perform variable selection means it is unable to exclude irrelevant variables from the model, thus hampering the interpretability of the model when it includes a large number of predictors. The fact that the mathematical relationships between inputs and outputs in non-linear models is often obscured, on the other hand, similarly prevents one from clearly interpreting how the included variables influence the response.

Table 3.1 includes two other families of methods—subset selection and regularisation. These use alternative model fitting procedures that can lead to improvements in both prediction accuracy and model interpretability (James et al., 2013). Critically, both of these approaches can perform variable selection. As we will see, however, the negative attributes of subset selection make these methods less desirable than regularisation.

The first paper in this thesis introduces several arguments for why regularisation techniques improve upon subset selection approaches in regression models. These arguments reference several publications that have highlighted the problems inherent to using stepwise selection methods for variable selection (Flom et al., 2007; Harrell, 2001; Hurvich and Tsai, 1990; Whittingham et al., 2006). In summary, the key issue is that subset selection approaches apply methods that are only appropriate for one test to many repeated statistical tests (Malek et al., 2007). In other words, methods for subset selection perform a large number of hypothesis tests, which violates the statistical theory of tests such as the F-test, which is based on one single hypothesis being tested (Harrell, 2001). The effects of this violation include an inflation in the probability of false positives (Type I errors) (Steyerberg et al., 1999; Mundry and Nunn, 2009; Wilkinson, 1979), an overestimation of the model's

R^2 value or fit (Cohen and Cohen, 1983), and standard errors, p values, and confidence intervals that are falsely low or narrow (Altman and Andersen, 1989). Furthermore, because subset selection is a discrete process wherein predictors are either kept in or removed from the model, small changes in the underlying data can lead to very different models being selected, presenting issues of model stability (Tibshirani, 1996). While these issues are relevant to both best subsets selection and stepwise selection, they are most often discussed in relation to the latter method. A more practical drawback to best subset selection is that it can be computationally infeasible as the number of predictors grows. Much more explanation and discussion of these issues is given in paper 1.

Regularisation methods avoid these violations of statistical theory because the selection of variables is based on the penalisation of coefficients (shrinkage) rather than on statistical testing (Desboulets, 2018). In other words, regularisation methods that have the benefit of sparsity (e.g., lasso and elastic net regression) do so not because they 'select' variables using iterative tests of significance but instead because they solve an optimisation problem wherein some coefficients are shrunk to zero, which is an indirect method for variable selection (Tibshirani, 1996). While subset selection methods often show an inflated Type I error rate due to multiple hypothesis testing, regularisation methods avoid this issue because variables are not selected using statistical tests. This fact does highlight one of the main negative attributes of regularisation, however, which is that the typical inferential constructs, such as standard errors and confidence intervals, cannot be easily computed for these methods. For this reason, developing approaches for inference and significance testing in such models is an active area of statistics research (Lockhart et al., 2014). Regarding problems related to statistical power, or the ability to detect Type II errors, there is mixed evidence on whether and how regularisation methods can avoid this issue, though there appears to be a trade-off between Type I and Type II errors (see Su et al. (2017)).

In addition to avoiding the various statistical issues mentioned above, regularisation has additional positive attributes that make it useful in high-dimensional situations and that enable it to handle multicollinearity while also preventing overfitting (Zou and Hastie, 2005).

These analysis issues present obstacles for modelling residential electricity consumption, and although regularisation methods have positive attributes that make them useful for handling these challenges, they are still rarely applied in energy research (Hsu, 2015). Paper 1 returns to this point to summarise and further explain these key analysis issues and how regularisation is well suited to addressing them.

Returning, then, to the thesis' methodological research question, there are several lines along which regularisation and, specifically, statistical learning with sparsity are expected to 'improve efforts' to characterise residential electricity consumption patterns. These have been briefly outlined in this section but are discussed in much greater detail in the individual papers where specific analysis issues are confronted and where regularisation methods are applied to address these. In addition, the discussion presented in Chapter 8 presents a summary of how and under what circumstances the methods chosen in this thesis are expected to improve models of residential electricity demand in order to provide guidance to future research.

Beyond the practical and theoretical justifications for using these methods in this thesis, the more general justification for using quantitative rather than qualitative methods throughout is two-fold: first, as this research aims to contribute to more detailed characterisations of demand, especially to highlight the key structural and occupant factors that affect consumption, it is sensible to use these types of methods given that they are more easily scaled, they can be applied to much different social, geographical, and cultural contexts, and they rely on data that are increasingly accessible to energy researchers.

There is no doubt, however, that other quantitative methods, such as engineering methods, are useful for a variety of purposes relevant to measuring and mapping electricity demand. The first paper provides a detailed discussion of these. Similarly, outside the realm of quantitative methods, there is a clear need for qualitative methods, such as ethnographic and interview-based approaches, to better understand complex human-energy dynamics and inform quantitative modelling approaches (Hiteva et al., 2018). To this end, there is a growing community of social scientific energy researchers who are providing exceptionally important insights for demand reduction and flexibility (Sovacool, 2014).

The second, and perhaps more important, reason that this thesis focuses on quantitative and, specifically, sparse statistical learning approaches to characterise demand is simply because of personal interest. There are many avenues one can take when starting a DPhil, and this is the research I decided would be most exciting and rewarding given my background and future professional interests.

3.5 Structure of analysis and research papers

The thesis draws on four empirical research papers that are principally differentiated by the temporal resolution at which they characterise electricity consumption patterns, but they also build upon one another in several other ways. This section first provides a general overview of the four papers. Next, it briefly summarises the papers to highlight how they build upon one another in terms of the methods applied, the empirical questions investigated, and the policy implications provided. As an extension to the infographic presented in Chapter 1, Table 3.2 presents the research questions specific to each of the four papers along with more detail on the statistical learning methods and evaluation procedures used in each. The publication status for each paper is given here, as well.

In general, the papers build upon one another in terms of their methodological approach and empirical scope. For this reason, the first paper is primarily a methods paper that introduces the concept of regularisation in statistical learning and shows how it can address many challenges inherent to energy demand research, using a case study of characterising annual electricity consumption patterns for a sample of California households. In papers 2 and 3, these methods are then applied to data sets of higher-resolution electricity and related household and time-use data. These three papers form the core methodological and empirical research of the thesis, though each has specific implications for policy, as well. The fourth paper is more explicitly focused on policy questions, though it also has methodological and empirical implications. It provides a framework for conducting time-

	Paper 1	**Paper 2**	**Paper 3**	**Paper 4**
Methodological research question	How can regularisation methods improve variable selection efforts and interpretability in statistical models of residential electricity consumption?	By how much does inclusion of time-use activity data in models of daily and intra-day electricity demand improve the explanatory power of these models?	What approaches can improve outcomes for clustering high-resolution residential electricity load profiles for the aim of identifying peak-period usage clusters?	How can a time-sensitive framework inform like-for-like comparisons of the savings potential of energy efficiency and flexibility measures in the U.S. building sector?
Empirical research question	Which structural and occupant factors are the strongest predictors of annual electricity consumption in California households?	How do the relationships between structural and occupant factors, including occupants' daily activities, and electricity demand vary as a function of the time of day in U.K. households?	What are the activity-based drivers of peak-period electricity demand patterns in U.K. households?	How do efficiency and flexibility measures compare in terms of their time-sensitive energy, cost, and emissions savings potential in the U.S. building sector?
Policy research question	How does identifying key drivers of annual residential electricity consumption inform energy efficiency policy design in California?	How can understanding which factors influence intra-day variations in household electricity consumption inform policies for demand flexibility in the U.K.?	How does evidence on the relationships between occupant activities and electricity consumption during peak hours provide insights for targeted interventions to reduce U.K. peak demand?	How does quantifying the time-sensitive impacts of U.S. building electricity use inform investments in national energy efficiency and flexibility technologies?
Statistical learning methods used	*Supervised learning (regression)* Lasso regression Ridge regression Elastic net regression Adaptive lasso regression Group lasso regression Stepwise regression *Resampling* Cross-validation *Other* Multiple imputation	*Supervised learning* Lasso regression OLS regression *Resampling* Cross-validation	*Supervised learning (classification)* Logistic regression with elastic net regularisation Random forests *Unsupervised learning* Cluster analysis *Resampling* Cross-validation Bootstrap resampling	N/A
Statistical learning model evaluation	Root mean-squared error (RMSE) R-squared Number of non-zero coefficients	Mean-squared error (MSE) R-squared Adjusted R-squared	Cluster validity indices (CVIs) Area under the Receiver Operating Characteristic (ROC) curve (AUC)	N/A
Publication status	Published in February 2019 in *Energy* after one round of peer review.	Published in April 2019 in *Energy Efficiency* after one round of peer review.	Published in December 2019 in *Applied Energy* after one round of peer review.	Published in December 2019 in *Environmental Research Letters* after two rounds of peer review.

Table 3.2: Overview of the four empirical papers submitted as part of this thesis.

sensitive analyses of electricity consumption and its resulting costs and GHG emissions, and it then quantifies the potential savings across these metrics for various efficiency and flexibility measures. More detail on each paper is provided next.

3.5.1 Paper 1

This paper argues that the use of regularisation methods in statistical models can address challenges inherent to building energy modelling while also enabling more accurate predictions and better identification of variables that influence consumption. It highlights that these approaches are seldom used in the building energy modelling literature.

Being primarily a study of different methods, this paper develops the intuition for using regularisation in studies of electricity consumption and then applies five such methods to original survey and electricity consumption data in a case study of the structural and occupant factors that influence demand in California households.

Its results show both which regularisation approaches perform the best in terms of yielding predictive yet sparse, interpretable models as well as which structural and occupant-related factors bear the strongest associations to annual electricity consumption patterns for the sample of households. The paper then discusses the implications of these findings for both improving building energy modelling as well as informing policies for demand reduction.

3.5.2 Paper 2

The second paper follows on from the first to show how the regularisation methods developed in the first paper can similarly improve model interpretability and highlight key factors associated with average daily electricity demand as well as intra-day variations in electricity demand over four-hour periods of the day. The paper was authored with my supervisor and a post-doctoral researcher for the METER project.

As its research question indicates, this paper similarly investigates structural and occupant factors, including many of the same socio-demographic and physical dwelling factors studied in the first paper, but it also incorporates more detailed appliance ownership data as well as categorised time-use data from occupants. The inclusion of time-use data is intended to bring occupant activities and their relationship to electricity demand at different times of day to the forefront of the analysis.

The data come from the METER project and include electricity data for 138 households. In this paper, one of the regularisation methods demonstrated in the first paper, lasso regression, is used to select key variables associated with daily and intra-day variations in electricity demand. Next, OLS regression is used to compare these variables and the strength of their relationship across four-hour periods of the day. An important component of this analysis includes estimating at what times of day and by how much does the inclusion of time-use data improve model explanatory power.

This paper provides insights for improving efforts to characterise electricity demand using time-use data as well as empirical findings on the types of activities and other household variables that have strong associations to electricity demand at different times of day. The findings have policy implications related to better understanding the constraints and opportunities for residential demand flexibility resulting from activity-led changes in demand.

3.5.3 Paper 3

The third paper uses the same data set as the second (albeit a larger sample size of 269 household electric load profiles) but takes a different statistical learning approach to characterise patterns of electricity demand during peak demand hours in the U.K. It has the same co-authors as the second paper.

Paper 3 applies cluster analysis, an unsupervised learning approach, to the task of characterising patterns of demand during peak hours (from 5 to 9 p.m.). Cluster analysis is one of the more common approaches used to characterise electricity consumption, especially hourly or sub-hourly usage data. But its use in this paper is essentially a pre-processing step to yield groups of households based on the similarities in their electric load profiles during evening peak hours. Where paper 3 builds upon the first two papers is in its subsequent application of regularisation techniques in a supervised classification setting to show how detailed occupant activity data can be used to predict whether a given household is more likely to follow a particular pattern of evening electricity demand.

As in paper 2, paper 3 is primarily concerned with understanding occupant activities and their relationship to electricity demand. It does so for the purpose of identifying the key activities that drive evening peak-time demand, which has implications for policies and strategies to reduce peak demand or shift electricity consumption to other times of day.

One final note on paper 3 that will be returned to in the discussion chapter: this paper is much more experimental in its design than the first two papers. Much of its formulation is motivated by an interest in contributing to the already vast literature on cluster analysis for segmenting electricity customers based on their load profiles, which is why it uses an unsupervised learning approach in the first part of the analysis. But it also aims to accurately predict evening peak-time demand clusters using detailed occupant activity data and regularisation techniques. The combination of these various approaches and methods builds on previous work and is novel in its inclusion of time-use data as well as its predictive modelling approach.

3.5.4 Paper 4

The final paper takes a different modelling approach than the previous three, as it does not directly employ statistical learning methods. This paper aims to provide an assessment of why a better understanding of the temporal aspects of residential electricity demand is important. In this way, the paper aims to answer the 'So what?' question for the thesis. It does this by developing a national baseline for time-sensitive analysis of the varying costs and emissions associated with electricity consumption in the U.S. building sector. This paper was co-authored by Jared Langevin, a Research Scientist at Berkeley Lab, who supervised my work while I was a research assistant there.

This paper is primarily focused on the policy relevance of attaining greater efficiency and flexibility in the building sector. The paper relies on a new methodological tool called Scout, a building efficiency impact assessment software developed by the U.S. DOE. It adapts this tool to show how building electricity demand and its resulting costs and emissions vary hourly by end use, region, and season in the U.S. In doing so, it creates the baselines necessary to enable a like-for-like comparison between static energy efficiency measures, which reduce electricity demand evenly across hours of the day, with dynamic energy flexibility measures, which shift and shed electricity use during specific hours of the day.

As new technologies enable buildings to more actively engage with the electric grid through smart appliances and controls, understanding how these can enable both a more efficient and more flexible building sector is increasingly important, especially in the context of an energy system that is rapidly transitioning to variable renewable power.

3.6 Chapter summary

This chapter has provided an introduction to statistical learning, the competing aims of prediction and inference, and the sub-fields of regularisation and sparsity. It has described some of the main statistical learning methods applied in this thesis and has given justifications for their use. The chapter also gave brief descriptions of the four empirical papers in order to show how they build upon one another in terms of the methods they use, their empirical scope, and their implications for energy efficiency and flexibility policy. Much more detailed methods sections are included in each of the papers, which are presented in the following four chapters.

3.7 References

Altman, D. G. and Andersen, P. K. (1989), 'Bootstrap investigation of the stability of a cox regression model', *Statistics in Medicine* **8**(7), 771–783.

Cohen, J. and Cohen, P. (1983), *Applied Multiple Regression/Correlation Analysis for the Behavioral Sciences*, 2nd ed edn, L. Erlbaum Associates, Hillsdale, N.J.

Desboulets, L. (2018), 'A review on variable selection in regression analysis', *Econometrics* **6**(4), 45.

Flom, P. L., Development, N., Institutes, R. and York, N. (2007), Stopping stepwise: Why stepwise and similar selection methods are bad, and what you should use, *in* 'NorthEast SAS Users Group Inc 20th Annual Conference', Baltimore, MD, p. 7.

Harrell, F. (2001), *Regression Modeling Strategies: With Applications to Linear Models, Logistic Regression, and Survival Analysis*, Springer Series in Statistics, Springer-Verlag, New York, NY.

Hastie, T., Tibshirani, R. and Wainwright, M. (2015), *Statistical Learning with Sparsity: The Lasso and Generalizations*, Chapman & Hall/CRC.

Hiteva, R., Ives, M., Weijnen, M. and Nikolic, I. (2018), A complementary understanding of residential energy demand, consumption and services, *in* C. Foulds and R. Robison, eds, 'Advancing Energy Policy: Lessons on the Integration of Social Sciences and Humanities', Springer International Publishing, Cham, pp. 111–127.

Hsu, D. (2015), 'Identifying key variables and interactions in statistical models of building energy consumption using regularization', *Energy* **83**, 144–155.

Hurvich, C. M. and Tsai, C.-L. (1990), 'The impact of model selection on inference in linear regression', *The American Statistician* **44**(3), 214–217.

James, G., Witten, D., Hastie, T. and Tibshirani, R. (2013), *An Introduction to Statistical Learning*, Vol. 103 of *Springer Texts in Statistics*, Springer New York, New York, NY.

Lockhart, R., Taylor, J., Tibshirani, R. J. and Tibshirani, R. (2014), 'A significance test for the lasso', *Annals of statistics* **42**(2), 413.

Malek, M. H., Berger, D. E. and Coburn, J. W. (2007), 'On the inappropriateness of stepwise regression analysis for model building and testing', *Eur J Appl Physiol* **101**(2), 263–264.

Mundry, R. and Nunn, C. L. (2009), 'Stepwise model fitting and statistical inference: Turning noise into signal pollution', *The American Naturalist* **173**(1), 119–123.

Osborne, J. W. and Waters, E. (2002), 'Four assumptions of multiple regression that researchers should always test', *Practical Assessment, Research & Evaluation* **8**(2), 1–5.

Shmueli, G. (2010), 'To explain or to predict?', *Statistical Science* **25**(3), 289–310.

Sovacool, B. K. (2014), 'What are we doing here? Analyzing fifteen years of energy scholarship and proposing a social science research agenda', *Energy Research & Social Science* **1**, 1–29.

Steyerberg, E. W., Eijkemans, M. J. C. and Habbema, J. D. F. (1999), 'Stepwise selection in small data sets: A simulation study of bias in logistic regression analysis', *Journal of Clinical Epidemiology* **52**(10), 935–942.

Su, W., Bogdan, M. and Candès, E. (2017), 'False discoveries occur early on the Lasso path', *Ann. Statist.* **45**(5), 2133–2150.

Tibshirani, R. (1996), 'Regression shrinkage and selection via the lasso', *Journal of the Royal Statistical Society. Series B (Methodological)* pp. 267–288.

Viegas, J. L., Vieira, S. M., Melício, R., Mendes, V. and Sousa, J. M. (2016), 'Classification of new electricity customers based on surveys and smart metering data', *Energy* **107**, 804–817.

Whittingham, M. J., Stephens, P. A., Bradbury, R. B. and Freckleton, R. P. (2006), 'Why do we still use stepwise modelling in ecology and behaviour?', *J Anim Ecol* **75**(5), 1182–1189.

Wilkinson, L. (1979), 'Tests of significance in stepwise regression', *Psychological Bulletin* **86**(1), 168–174.

Zou, H. and Hastie, T. (2005), 'Regularization and variable selection via the elastic net', *Journal of the Royal Statistical Society. Series B (Statistical Methodology)* **67**(2), 301–320.

Chapter 4

Investigating structural and occupant drivers of annual residential electricity consumption using regularisation in regression models (Paper 1)

Aven Satre-Meloy
Published in *Energy*

Abstract

Achieving further reductions in building electricity usage and emissions requires a detailed characterisation of electricity consumption in homes. Understanding drivers of consumption can inform strategies for promoting conservation and efficiency. While there exist numerous approaches for modelling building energy demand and identifying key drivers, the use of regularisation methods in statistical models can address challenges inherent to building energy modelling while also enabling more accurate predictions and better identification of variables that influence consumption.

This paper applies five regularisation techniques to regression models of original survey and annual electricity consumption data for more than one thousand households in Palo Alto, California. It finds that of these, elastic net and two extensions of the lasso—group lasso and adaptive lasso—outperform other approaches in terms of prediction accuracy and model interpretability. These findings contribute to methodological approaches for modelling energy consumption in buildings as well as to our understanding of key drivers of consumption. The paper shows that while structural factors predominate in explaining annual electricity consumption patterns, habitual actions taken to save energy in the home are important for reducing consumption while pro-environmental attitudes and energy literacy are not. The implications of

these findings, both for improving building energy modelling and for informing demand reduction strategies, are discussed in the context of the low-carbon transition.

4.1 Introduction

The U.S. is the second largest energy market and emitter of greenhouse gases (GHG) in the world after China, and buildings hold the largest share of U.S. energy consumption at 41%, more than half of which comes from the residential sector (Nejat et al., 2015). Residential energy demand has remained relatively stable since 1990, yet the sector still accounted for 20% of CO_2 emissions from fossil fuel combustion in 2015. Of these, 68% were attributable to electricity consumption for lighting, heating, cooling, and operating appliances, with the remainder due to consumption of other fuels for heating and cooking (U.S. EPA, 2017). Electricity demand is projected to grow over the next thirty years, in part because of an increase in the adoption of cooling technologies due to future warming (U.S. EIA, 2018).

Demand reduction is believed to play an important role in the effort to reduce emissions from the residential building sector. The EPA predicts that demand-side efficiency and saving measures could result in a net cumulative demand reduction of 7.83% by 2030 (U.S. EPA, 2015). These measures are also some of the most cost effective for reducing emissions from the power sector, delivering savings at a fraction of the retail cost of electricity (Hoffman et al., 2018). Only recently have demand-side strategies begun to receive closer scrutiny in national and global scenarios for limiting warming to the 1.5°C target agreed in Paris (Grubler et al., 2018). This emerging literature stresses the importance of demand-side measures for achieving ambitious climate goals and delivering societal co-benefits for health, equity, and security (Mundaca et al., 2018).

The residential building sector's large share of electricity consumption and sizeable potential for reducing emissions warrant detailed investigations into the drivers of consumption. A deeper understanding of the characteristics of electricity consumption in homes can inform strategies and policies for promoting conservation and efficiency. This is especially important given that much of the existing quantitative research on building energy consumption and prediction has focused on non-residential buildings (Amasyali and El-Gohary, 2018).

Approaches to investigating drivers of building electricity consumption have proliferated in recent years alongside a similar expansion in available data to analyse these drivers. Both statistical and engineering techniques are increasingly applied to diverse, multivariate data to quantify the contribution of different factors to household electricity consumption. These methods benefit from improved computing power, access to large data sets, and new algorithmic approaches for modelling electricity consumption.

Yet despite their proliferation, statistical and engineering methods for modelling building energy consumption face numerous challenges, especially in the context of informing policy development. Hsu (2015) summarises key challenges that are shared across energy analysis research, and several of these are highlighted here.

First the number of factors that possibly influence energy consumption, including structural factors, such as physical dwelling characteristics and efficiency standards, as well as economic, social, and behavioural dimensions, is almost limitless. Understanding the comparative contributions of these different factors to consumption patterns can improve intervention efforts to promote conservation. Second, although the availability of data is improving, it is still difficult and expensive to gather comprehensive data on these factors, so results are often based on small data sets specific to particular geographic, economic, and social contexts. Third, statistical models based on small samples often do not have high out-of-sample predictive accuracy. Especially when the set of possible predictive factors is large (and in 'high-dimensional' problems, larger than the number of observations), models often 'overfit' the data, meaning they do not generalise well to new data and lead to poor predictions and inferences. Fourth, including a large number of predictors in statistical models increases the likelihood of multicollinearity, where multiple predictors have high degrees of pair-wise correlation, which can inflate the standard errors of coefficients in statistical models and lead to misinterpretation (Swan and Ugursal, 2009). Finally, an additional challenge is the prevalence of missing data, which is common in data sets pulled together from numerous sources, especially from household surveys where completion is not mandatory. Missing data, if not handled properly, can result in loss of information and introduce bias (Pampaka et al., 2016).

Overcoming these analytical challenges is important for interpreting model results accurately and properly informing strategies for delivering energy savings, but many of these issues are not well addressed in the energy consumption literature, and statistical techniques to handle these challenges are rarely applied in empirical energy consumption studies (Hsu, 2015). As the following review of literature will show, numerous modelling techniques exist to estimate residential energy consumption, but many of these are geared toward improving predictive performance without also yielding interpretable results. This phenomenon has become increasingly common with advanced machine learning approaches, especially those in the field of deep learning. While improvements in prediction are certainly important for numerous purposes, developing better solutions for reducing energy demand requires interpretable models that identify important factors explaining consumption. Thus, statistical approaches that can improve predictive performance while also ensuring a more robust variable selection process are especially relevant for residential electricity consumption research.

This paper therefore makes two primary contributions. First, it contributes to the literature on model selection for electricity consumption by applying regularisation techniques to linear regression models of annual electricity consumption. Following Hsu's (2015) introduction of these techniques to the energy consumption literature several years ago, they continue to be seldom-used despite their demonstrable benefits for improving statistical models and identifying key variables. This paper will show how the use of regularisation techniques should be guided by the analysis objective and the structure of the data. It shows how several recent extensions to these methods can improve results for prediction and interpretation objectives when the data contain many different types of variables, which is common in residential energy demand research. The second aim of this paper is empirical, demonstrating the use of these techniques on an original data set of annual electricity usage data and a wide range of structural and occupant factors for over 1,000 households in California.

The paper is organised as follows: Section 4.2 reviews related work, both on statistical modelling of energy consumption in buildings as well as on determinants of consumption. It highlights areas of uncertainty and gaps in our knowledge. Section 4.3 describes the use of regularisation methods, including several recent extensions, and the statistical motivations for the modelling approach undertaken in this paper. Section 4.4 describes data collection, organisation, and preprocessing procedures. Section 4.5 presents results. Implications for both modelling and policy are discussed in Section 4.6, and Section 4.7 concludes with a discussion of how the methods used in this paper can inform further research in building energy consumption analysis.

4.2 Related work

This review of related work is split into two sections. Section 4.2.1 describes the high-level taxonomy of approaches for building energy consumption modelling and then provides a more detailed review of statistical methods and several key issues that are present, including the competing aims of prediction and interpretation and the need for robust variable selection techniques. Section 4.2.2 reviews the literature on determinants of household electricity consumption.

4.2.1 Approaches for modelling building energy consumption

Swan and Ugursal (2009) review residential energy consumption models and show that several approaches are appropriate, depending on the scale of interest. These approaches are either top-down or bottom-up, and Figure 4.1 shows the methods common to each. Top-down models use large, statistical databases to quantify regional or national energy supply requirements. Econometric models are defined by their inputs, which are typically macroeconomic indicators, whereas technological models generally use characteristics of the entire housing stock, such as appliance ownership. These models are useful for predicting trends in consumption for national planning purposes, but they require little detail beyond these broad indicators and thus provide limited insight into the micro-scale factors that influence consumption, including occupant behaviour.

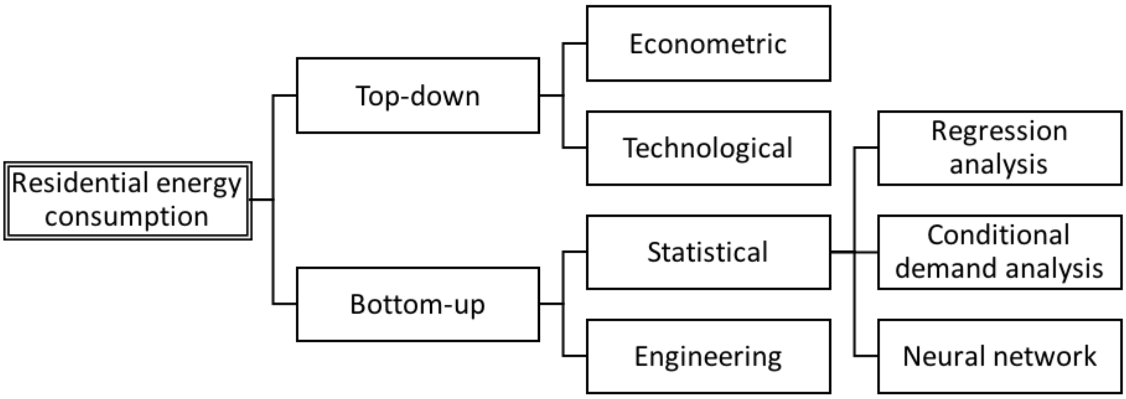

Figure 4.1: Modelling techniques for estimating residential energy consumption. Adapted from Swan and Ugursal (2009).

Bottom-up models, on the other hand, account for energy consumption due to individual end uses and can use a variety of input data. These data can include socio-demographic, occupant behaviour, or technology factors. There are two distinct categories of bottom-up models, which use different approaches for estimating consumption. Engineering methods, also called building physics models, use detailed data on dwelling characteristics, power ratings and use of appliances, and thermodynamic principles to predict consumption. Statistical methods instead use mathematical principles to describe the relationship between predictive variables and household electricity consumption.

The benefits of engineering methods include the use of physically measurable data to determine the consumption of specific end uses and technologies. Measurements and simulations are useful for describing existing technologies in greater detail and modelling the prospective impact of new technologies. The drawback of using these models is that they rely on assumptions about occupant behaviour, do not include other socio-demographic or economic data, and usually require a lot of technical data and measurements of building characteristics while requiring more computational power for analysis (Kavgic et al., 2010).

Statistical methods, on the other hand, can incorporate more varied socio-demographic and behavioural data, are often less computationally intensive, and are somewhat easier to develop and use. Several exceptions include non-linear models, which are discussed in greater detail below.Because statistical models represent a purely mathematical relationship between energy consumption and predictive variables, however, they are often prone to more error and uncertainty than engineering models (Swan and Ugursal, 2009; Zhao and Magoulès, 2012). Given recent advances in statistical modelling, and given that statistical modelling techniques are employed in this paper, a brief review of these is provided in the following section.

Statistical and data-driven models

The main approaches for statistical modelling highlighted in Swan and Ugursal (2009) are regression analysis, conditional demand analysis (CDA), and artificial neural networks (ANN). More recent reviews include additional methods such as support vector machines (SVM) and decision trees (DT) (Amasyali and El-Gohary, 2018; Wei et al., 2018; Zhao and Magoulès, 2010). Each of these are briefly described in turn. For a more complete review of these methods and their mathematical properties, see Wei et al. (2018).

Regression analysis is one of the most common approaches for modelling building energy consumption. In its simplest form, regression analysis determines the size and direction of associations between predictive factors and electricity consumption. Predictors are selected based on expectations of what drives consumption and data that are available or collected. Selecting predictors is the subject of a broad statistical literature, which is further discussed in Section 4.2.1. Models are evaluated using goodness-of-fit measures and model predictive error. Key predictors and their coefficients are examined to determine the strength and statistical significance of their relationships with consumption. Regression models are simple to develop and use, yet they require access to large sets of historical data and do not often achieve the predictive accuracy of other methods.

Conditional demand analysis (CDA) uses regression analysis but only includes as predictors the various end-use appliances owned in the dwelling. The coefficients in the model thus represent the use level and rating of the appliances. While this technique is relatively simple to use, it requires detailed data on household appliance ownership and a large sample of dwellings (Swan and Ugursal, 2009).

Artificial neural networks (ANNs) have been used for several decades and have seen a resurgence with the rise of machine learning disciplines and, especially, deep learning approaches. The method is based on analytic techniques originally developed for studying human neurophysiology. The simplest ANNs include three layers: an input layer, a hidden layer, and an output layer, each of which has interconnected neurons that send signals to the neurons in sequential layers using an activation function (Amasyali and El-Gohary, 2018). The reason ANNs have gained such popularity is their ability to model incredibly complex, non-linear relationships. The trade-off for this gain in model complexity is that the coefficients in the model do not have physical significance, so interpreting the influence of different factors in neural networks is challenging. In addition, ANNs are prone to overfitting and often have poor out-of-sample predictive accuracy.

Another popular method in machine learning is the SVM, which also performs particularly well when the relationship between the inputs and the response is non-linear. Support Vector Regression (SVR) is the application of SVM principles to regression problems. SVR works by mapping data inputs to a higher dimensional feature space using a kernel function and then constructing a linear model that keeps the error within a predefined threshold. It has shown improved predictive capabilities for building energy consumption (Zhao and Magoulès, 2012). An additional benefit of SVMs is that they require fewer parameters and less training data. Like ANNs, however, SVMs are more complex models that suffer from computational inefficiencies, though optimisation of these algorithms is the subject of research (e.g., Zhao and Magoulès, 2010).

Decision trees (DT) work by partitioning data into groups based on predefined predictor variables, where each variable represents a root or branch in the tree, and the data are partitioned into smaller groups along the branches. The modeller chooses which variables to use as nodes and can decide where to trim the DT. In this way, a DT visually represents the data partitioning decisions made at each branch, and for this reason they are simple to understand and interpret, which is one of their main advantages. They have also proven effective in building energy prediction (Yu et al., 2010). With larger numbers of predictors, DTs can become overly complex, but ensemble methods such as random forests can help prevent overfitting (Wang et al., 2018).

These methods are some of the most common for statistical modelling of energy consumption, but there are many others and many variations of each of these. One of the key differences between these methods, however, is whether they are used primarily for prediction of building energy consumption or explanation of factors that influence consumption. Much of the research interest in machine learning methods such as ANNs, SVMs, and DTs is improving the ability to predict consumption. While this is certainly important in building energy research, the pursuit of more accurate predictions can hamper interpretation efforts. Increasing gains in model accuracy often relies on increased model complexity, which is commonly the case for ANNs and SVMs. This complexity makes it difficult to understand the relationships between data inputs and the response. While linear regression models may not match the predictive accuracy of complex, non-linear models, they are interpretable and can clearly describe relationships between variables and energy consumption. When this is the aim, model simplicity is essential. This gives the model practical significance for informing strategies to reduce demand.

For the interested reader, an insightful essay comparing the objectives of prediction and explanation in statistical models is given by Shmueli (2010). The essay concludes that while these objectives often delineate the choice of variables, methods, and approaches for selecting, validating, and evaluating statistical models, in most cases it is appropriate to consider both the predictive and explanatory power of models. Even when the objective is not primarily prediction, the predictive qualities of a model should be reported in research, and *vice versa*. Model performance can then be judged based on both of these criteria.

Variable selection and related challenges in statistical models

Variable selection is a key issue in statistical modelling, especially when the number of candidate variables is large. For data with p variables, the number of possible models with subsets of the p variables is 2^p. A seemingly simple data set with 10 variables gives over 1,000 possible models with subsets of variables. Even with modern computing capabilities, constructing every possible model and comparing each using some evaluative criteria is impractical as the number of candidate variables grows.

Two additional challenges are more likely to be present when the number of candidate variables increases. The first is multicollinearity between predictors. Multicollinearity exists when one predictor variable has a near linear relationship with another (Farrar and Glauber, 1967). When this is the case, the coefficient estimates for the regressors become unstable and are susceptible to erratic changes with small changes to the model or data. Multicollinearity has been identified as a challenge in energy consumption modelling in numerous instances (Fumo and Rafe Biswas, 2015; Huebner et al., 2016; Swan and Ugursal, 2009). It is a challenge unique to the objective of constructing explanatory models and interpreting variable size and significance, as predictive accuracy does not suffer when multicollinearity effects are present (Shmueli, 2010).

The second challenge for models with many potential predictors is overfitting. Overfitting occurs when the model is overly complex or includes more variables than necessary. In high-dimensional cases, where models have more predictors (p) than observations (n), approaches such as OLS regression do not have well-defined solutions. The result of overfitting is a model that fits so well to the existing data that it does not generalise to new data. When models are overfit, they suffer from high variance, meaning they capture the noise inherent in the data along with the underlying patterns. High bias models, on the other hand, are too simple and do not fit well to the existing data. There is a well-researched trade-off between bias and variance in the statistics literature (Friedman et al., 2001). In the energy modelling literature, efforts to address overfitting are most common in predictive modelling or forecasting studies (e.g., Jain et al., 2014; Pao, 2009).

These challenges stand out in efforts to model energy consumption, especially for explanatory purposes, because of the sheer number of potential factors that influence usage and because of their potential for high pairwise correlation. Certainly, domain knowledge and previous research should guide the selection of relevant variables, but analyses that explore large sets of untested variables are also valuable, and statistical techniques that can address the challenges of large predictor sets, multicollinearity, and overfitting can aid in variable selection efforts. For this reason, there is a rich and active literature in statistics on variable selection (Heinze et al., 2018).

Some of the most popular statistical techniques for selecting variables fall under the stepwise family of approaches, which includes forward selection, backward elimination, and stepwise selection (Heinze et al., 2018). These procedures iteratively construct regression models by adding or removing predictors based on a test statistic or minimising an evaluative criterion, such as the Akaike information criterion (AIC) or Bayesian information criterion (BIC), until a final model is attained. Stepwise regression techniques have been applied in numerous studies of energy consumption to identify relevant predictors (Deb and Lee, 2018; Filippín et al., 2013; Kialashaki and Reisel, 2013; Wang et al., 2016). Other approaches for variable selection in energy consumption studies include principal components regression (PCR) and partial least squares regression (PLSR) (Djuric and Novakovic, 2012; Lam et al., 2008; Ndiaye and Gabriel, 2011).

Stepwise regression as an approach to variable selection has been derided in the statistics literature for violating statistical theory and causing important practical consequences for analysis (Harrell, 2001; Ratner, 2010). Some of these issues include R-squared values and regression coefficients that are biased on the high side, severe problems handling multicollinearity, and predicted values that are falsely narrow. PCR and PLSR do not have these same issues but do present challenges for interpretation because they transform predictor variables into linear combinations of the original predictor variables.

A separate class of variable selection techniques that can address many of these issues is regularisation. Regularisation methods, also called penalised regression methods, have received substantial attention in statistical research (Friedman et al., 2001), but their application to statistical modelling of energy consumption is still surprisingly rare. When Hsu (2015) first showed how the application of regularisation methods could improve efforts to identify key factors influencing consumption, the author's review of three prominent energy journals (*Energy*, *Energy Policy*, and *Applied Energy*) showed only a few papers applying these methods, mostly in econometric analyses. An updated search in these journals confirms they continue to be seldom used. Fewer than a total of 20 papers in these journals (including *Energy and Buildings*) use regularisation methods in modelling energy consumption, and much of their use is concentrated in recent machine learning analyses (Deng et al., 2018; Robinson et al., 2017) or in energy forecasting studies (Guo et al., 2018; Jain R. K. et al., 2014; Suryanarayana et al., 2018; Uniejewski et al., 2016). In two cases, these techniques have been used to analyse drivers of residential energy consumption in the U.K. and France (Belaïd et al., 2019; Huebner et al., 2016).

Regularisation methods are primarily used to prevent overfitting, but in some cases they are appropriate for handling multicollinearity and also variable selection. They also have consistently shown improved predictive ability in statistical models because they sacrifice some model bias for a sizeable reduction in the variance of predicted values. A full description of these methods and several recent extensions is given in Section 4.3.1.

4.2.2 Determinants of residential electricity consumption

While the previous section provided a review of related work in the energy modelling literature, this section will provide a review of literature investigating determinants of residential electricity consumption.

Jones et al. (2015) provided the first systematic review of international research investigating the determinants of electricity consumption and found that at least 62 factors have been studied, but only 20 of these were shown to unambiguously and consistently show a significant positive effect on electricity use. The authors found that the number of papers confirming a positive effect on consumption is much higher than the number showing a significant negative effect. Factors considered in the literature include: socio-demographic factors, physical dwelling characteristics and appliance ownership, occupant attitudinal factors and energy literacy, and occupant behaviour. Each of these factors will be reviewed in turn in the following sections.

Socio-demographic factors

Of the many possible occupant socio-demographic indicators to investigate, most studies focus on gender, age, and number of occupants, household income, and tenure of the dwelling (whether it is owned or rented). Nearly all studies reviewed show that the number of occupants has a significant, positive effect on household electricity consumption (e.g., Bedir et al., 2013; Ndiaye and Gabriel, 2011). The presence of young adolescents tends to amplify this trend (Brounen et al., 2012; Wyatt, 2013). Wiesmann et al. (2011) show that per capita electricity consumption is lower in households with more occupants, and Kavousian et al. (2013) find that the rate of usage increase slows with every doubling in occupancy.

The gender of the homeowner is not often statistically significant in regression models for household electricity usage, though Brounen et al. (2012) find per capita usage to be lower in dwellings occupied by females even after controlling for wealth.

Age of the occupants shows conflicting associations to usage. Several studies find a negative correlation between age and consumption (Brounen et al., 2012; Kavousian et al., 2013; McLoughlin et al., 2012), while others find a positive correlation (Guerra-Santin and Itard, 2010; Huebner et al., 2016; Tiwari, 2000). Researchers attribute these disparities to the fact that, in some cases, older occupants tend to be more aware of their consumption and use fewer electronic gadgets, but, in others, they spend more time in the home and are thus likely to consume more electricity.

Two other socio-demographic indicators that have often shown significant effects on household electricity consumption are income and tenure. The results on household income are also mixed: numerous studies find a monotonic and positive relationship between household income and electricity consumption (Bedir et al., 2013; Brounen et al., 2012; Cayla et al., 2011; Cramer et al., 1985; Kelly, 2011; Yohanis et al., 2008), but others find that the effect is small when controlling for other variables (Bedir et al., 2013; Kavousian et al., 2013; Wiesmann et al., 2011).

Home ownership is associated with higher electricity usage in (Wyatt, 2013) and (Wiesmann et al., 2011), but it shows no significant relationship in (Kavousian et al., 2013).

Physical dwelling characteristics

The size of a dwelling explains a large percentage of the variance in consumption (Brounen et al., 2012; Gram-Hanssen, 2013; Huebner et al., 2016; Kavousian et al., 2013; Wyatt, 2013), with detached dwellings using more electricity than apartments or flats (Huebner et al., 2016; Kavousian et al., 2013; Wyatt, 2013; Yohanis et al., 2008).

Older houses are shown to consume more electricity, likely due to less efficient building fabrics (Lindén et al., 2006; Wyatt, 2013), but some studies do not find this effect statistically significant (Brounen et al., 2012; Kavousian et al., 2013).

Even efficiency measures, such as insulation or double-glazed windows, are shown to have mixed relationships with consumption. Some studies find that they do reduce usage (Assimakopoulos, 1992; Hirst and Goeltz, 1985; Kavousian et al., 2013); others find no correlation (Cramer et al., 1985) or even a positive correlation (Kelly, 2011). One explanation given is that insulation measures are often correlated with house size and income.

Ownership of air conditioning (AC) significantly and consistently increases electricity usage (Cramer et al., 1985; Ndiaye and Gabriel, 2011; Tso and Guan, 2014; Tso and Yau, 2007), more so for central AC than window units. Results are sensitive to the climatic conditions where the study took place (Sanquist et al., 2012).

Ownership of more appliances generally correlates to greater electricity consumption (Bedir et al., 2013; Huebner et al., 2016; Jones et al., 2015; Wiesmann et al., 2011).

Ownership of devices that are intended to save electricity, including programmable and smart thermostats, smart meters and in-home displays, LED lighting, and others, are not as often included in empirical studies. The role of feedback and its affect on consumption is an area of growing interest (Abrahamse et al., 2007; Darby, 2006; Fischer, 2008; Schultz et al., 2015; Van Houwelingen and Van Raaij, 1989). These studies suggest potentially significant savings.[1]

Electric vehicles (EV) are a new class of electricity use and can lead to significant increases in household electricity consumption (Huang et al., 2012). U.S. DOE (2014) find that ownership of some EV models can double the electricity consumption of a single-family home.

Occupant attitudinal factors and energy literacy

The literature includes occupant attitudes on care for the environment, concern for climate change, and support for energy conservation and renewable energy. Energy literacy, or the extent to which individuals are familiar with and understand key concepts and issues related to energy, and its relationship with electricity usage has not been studied extensively.

Several studies that measure pro-environmental attitudes by asking respondents to rate their level of agreement with environmental statements find that attitudes cannot explain historical electricity consumption patterns but can explain savings in intervention studies or the occupant's self-reported engagement in energy-saving behaviour (Abrahamse and Steg, 2009; Brandon and Lewis, 1999; Gadenne et al., 2011). Vringer et al. (2007) find no significant differences in consumption for groups of households with different value patterns, and Bartiaux and Gram-Hanssen (2005) conclude that it would generally be difficult to use attitudes toward the environment to explain differences in electricity consumption between

[1]See Ehrhardt-Martinez et al. (2010) for a meta-review of 36 energy feedback studies from 1995–2010.

countries.

In the few studies where it is included, energy literacy is not found to correlate significantly with either historical consumption or energy conservation behaviour (Brounen et al., 2013). The National Environmental Education & Training Foundation (NEETF) gave a short energy knowledge quiz to a nationally representative sample of 1,503 Americans to determine the public's basic knowledge of energy issues. NEETF's report claims that "higher levels of knowledge of energy production, consumption, and conservation... have a positive effect on the likelihood of engaging in day-to-day activities that directly or indirectly conserve energy or benefit the environment" (NEETF & RoperASW, 2002, p. v). However, the actual reduction in demand was not measured, leaving a gap in our knowledge of the potential of more energy-informed citizens to reduce demand. Only 12% of Americans passed a basic quiz on energy topics, even though 75% rated themselves as having either 'a lot' or 'a fair amount' of knowledge about energy. Energy literacy has likely not been included in empirical studies as often as other factors because it is inherently difficult and subjective to measure. Most studies that measure energy literacy rates, including the NEETF study, do so with quizzes that ask questions about how and where energy is generated and consumed.

Occupant behavioural factors

Occupant behaviours influence electricity usage (Gram-Hanssen, 2014), and some studies investigating this relationship conclude that usage reductions of 10–20% are achievable by modifying behaviours alone (Mansouri et al., 1996).

Studies of conservation behaviour generally examine either 'habitual' actions or 'purchasing' activities (Barr et al., 2005). Gardner and Stern (2009) distinguish between these by specifying the former as 'curtailment' behaviours and the latter as 'efficiency' behaviours. They suggest that efficiency-improving actions yield greater savings than curtailing the use of appliances, lights, or inefficient equipment. The other main difference between the two is that curtailment actions must be repeated continuously over time, whereas efficiency measures need only be taken once or a few times and do not require continuing attention and effort. The authors' list of the most effective behaviours inside the home includes turning down the thermostat during the night and curtailing AC use during the day.

A number of studies find that occupant behaviours are important in explaining usage when controlling for structural elements (Brohus et al., 2010; Steemers and Yun, 2009; Thøgersen and Grønhøj, 2010; Wallis et al., 2016). Huebner et al. (2013) warn that similar findings from their study may not be generalisable.

Long-term curtailment behaviour is also measured in Kavousian et al.'s (2013) research with inconclusive findings on its impact. They find that, contrary to their expectations, the behaviour of 'Purchasing Energy-Star Appliances and Air Conditioners' is positively associated with households' daily minimum electricity consumption. They offer, as a possible explanation, the much-studied rebound effect where an increase in appliance or device efficiency results in increased use of them (Sorrell et al., 2009). They also find that those who report a long-term habit of 'Turning Off Lights When Not in Use' consume more electricity on average. This gap between individuals' intentions is investigated by Kennedy et al. (2009), who find that 72% of respondents self-reported a gap between their intentions and their actions related to the environment.

4.3 Methodology

The above literature review highlights a notable gap in the use of regularisation methods for energy use models and inconclusive findings on determinants of consumption. The challenge of variable selection looms large in studies of residential electricity usage, and related analytical challenges present difficulties for model interpretation. This section describes the fundamental regularisation methods and their application to multiple linear regression models. It first introduces these methods and then describes the motivations for using regularisation in this study. It then describes two recent extensions and how they improve on some of the shortcomings of the original regularisation methods. Next, it describes the model training, testing, and validation procedures. Lastly, it describes an important step taken during data preprocessing to address the issue of missing data.

4.3.1 Regularisation: overview and motivation

Regularisation methods are known as shrinkage methods because they shrink the coefficients of regression predictors, which trades off a small increase in model bias for a greater reduction in variance. The methods do this by applying a penalty term to the least squares estimator, hence the name 'penalised regression'. In the typical regression situation, we have data (x_i, y_i), $i = 1, 2, ..., n$, where x_i and y_i are the regressors and response for the ith observation, respectively, and where x_j denotes the jth predictor, $j = 1, 2, ..., p$. In OLS regression, we aim to estimate predictor coefficients (β_j) by minimising the residual sum-of-squares with respect to β:

$$RSS(\beta) = \sum_{i=1}^{n} \left(y_i - \beta_0 - \sum_{j=1}^{p} \beta_j x_{ij} \right)^2 \tag{4.1}$$

Penalised regression methods constrain this optimisation problem by adding a penalty term in the estimation of model coefficients. Because the penalties depend on the magnitude of these coefficients, the predictors and response are centered and standardised to have mean zero and a standard deviation of 1.

The three fundamental methods in regularisation are ridge regression, developed by Hoerl and Kennard (1970), lasso regression, introduced by Tibshirani (1996), and the elastic net, introduced by Zou and Hastie (2005). The penalty term in each case is slightly different. In ridge regression, the penalised optimisation problem is:

$$\hat{\beta}^{ridge} = \underset{\beta}{\text{argmin}} \left(RSS(\beta) + \lambda \sum_{j=1}^{p} \beta_j^2 \right) \tag{4.2}$$

where $\lambda \geq 0$ is the parameter that controls the amount of shrinkage. As λ increases, so does the penalty, which in ridge regression is the sum-of-squares of the coefficients. For this reason, ridge regression is also called l_2-regularisation because it constrains coefficients by their l_2 norm. Penalising by $\sum_{j=1}^{p} \beta_j^2$ has the effect of shrinking model coefficients but never to zero. An interest in yielding sparse, interpretable models is what motivated the introduction of the lasso, which constrains coefficients by their l_1 norm, and is given by:

$$\hat{\beta}^{lasso} = \underset{\beta}{\operatorname{argmin}} \left(RSS(\beta) + \lambda \sum_{j=1}^{p} \mid \beta_j \mid \right) \tag{4.3}$$

In lasso regression, this penalty constraint delivers sparsity, meaning some coefficients are set exactly to zero. In this way, beyond improving prediction, lasso performs variable selection and thus provides a level of interpretability in the model.

The motivation for the third regularisation method is due to the behaviour of lasso given highly correlated predictors. In lasso regression, the penalty tends to set only one of the predictor's coefficients to zero, and this procedure can yield non-unique solutions as well as poorer predictions when important but correlated predictors are removed from the model. A combination of both ridge and lasso penalties is the elastic net penalty:

$$\lambda \sum_{j=1}^{p} (\alpha \beta_j^2 + (1 - \alpha) \mid \beta_j \mid) \tag{4.4}$$

where α is an additional tuning parameter that can be tuned to constrain the optimisation by both the l_1 and l_2 norms. Eq. 4.4 is a generalised formulation of the three regularisation penalties. If $\alpha = 1$, this penalty is the ridge penalty, and if $\alpha = 0$, it is the lasso penalty.

These three methods have properties that make them useful in different situations, so it is important that their use is guided by the objective of the analysis. Several points on the motivations for using regularisation in this study are thus provided here.

While all three methods are able to reduce model variance and prevent overfitting, the most important difference between the three is whether or not they give a sparse solution. Both lasso and the elastic net penalties can yield sparsity in the model, whereas ridge cannot.

In addition, because ridge regression penalises coefficients by adding the sum-of-squares of the coefficients, it penalises the largest βs more than it does the smaller ones. This is sometimes important when inspecting ridge solutions, as it may be difficult to interpret which predictors are influential in the model.

As was mentioned previously, the behaviour of the regularisation methods when multicollinearity is present is somewhat different. The elastic net is said to have a *grouping effect*, which follows the intuition that highly correlated predictors will likely have similar estimated coefficients, and by combining ridge and lasso penalties, it keeps groups of correlated predictors in the model (Zou and Hastie, 2005). Ridge regression shrinks highly correlated predictors' coefficients toward one another, but this effect is still preferable to the situation with lasso, which tends to arbitrarily set one of the highly correlated predictors to zero. Elastic net is thus the preferred method when both multicollinearity is present and sparsity is an objective of the analysis. When multicollinearity is not an issue and pairwise correlations between predictors are low, there is often little difference in predictive accuracy between lasso and elastic net.

In high-dimensional situations where $p \gg n$, lasso can at most select n predictors, which was shown by Zou and Hastie (2005) to be a limiting feature in variable selection. In these situations, both elastic net and ridge can select more than n predictors and are preferable for more accurate models. Again, the elastic net is preferred over ridge in situations where sparsity is a goal.

The motivations for using regularisation in this paper are thus guided by the following characteristics of the data. First, the number of predictors ($p = 60$) is not larger than the number of observations ($n = 1008$), so the data are not high-dimensional even though the predictor set is quite large.

Second, multicollinearity among predictors does not appear to be an issue. Multicollinearity can be investigated by inspecting variance-inflation factors (VIFs), which signal whether regression coefficients are inflated due to correlation between predictor variables; if they are uncorrelated, VIF = 1. Traditionally, VIFs greater than 10 indicate high multicollinearity (Roberts and Thatcher, 2009), but recent work suggests that the cut-off point for VIFs should be much lower—Diamantopoulos (2006) set the limit at 3.3. The VIFs for the predictors in the present data range from 1.08—2.44 with a mean of 1.48, well below the cut-off point that signals potential issues.

This study has the stated aims of addressing overfitting to improve predictive performance while also performing variable selection to construct a more parsimonious model for interpretation purposes. The data are not high-dimensional (though still include more predictors than would be tractable using best-subsets methods), and pair-wise correlations between predictors are low. For this reason, lasso and elastic net are expected to perform similarly. As the next section will show, however, two extensions of the lasso may be expected to improve on both aims stated in this paper, given several additional aspects of the data.

4.3.2 Extensions of the lasso: group and adaptive lasso

The previous section discussed some of the situations where lasso regression does not perform adequately, such as when multicollinearity effects are present or in $p \gg n$ situations. An additional challenge for lasso regression is handling categorical predictors. Yuan and Lin (2006) showed that lasso is designed to select individual predictors rather than groups of predictors. As categorical predictors are normally coded as multiple dummy variables, where each dummy represents a different category, it makes sense in analysis to consider these variables together rather than separately when applying the shrinkage penalty. It would be inappropriate to include some of these variables in the model but not others. Especially in the building energy domain, categorical predictors are quite common (e.g., type of building, type of heating system, ownership structure).

To address this drawback of the lasso, Yuan and Lin (2006) introduced the group lasso (*gLasso*), a generalisation of the standard lasso optimisation problem. With p predictors divided into L groups, where p_l is the number in group l, and where, for ease of notation, X_l represents the predictors corresponding to the lth group, with corresponding coefficients

β_l, the group lasso solves the convex optimisation problem:

$$\hat{\beta}^{gLasso} = \underset{\beta}{\operatorname{argmin}} \left(\|y - \sum_{l=1}^{L} \beta_l X_l)\|_2^2 + \lambda \sum_{l=1}^{L} \sqrt{p_l} \|\beta_l\|_2 \right) \tag{4.5}$$

where$\sqrt{p_l}$ accounts for the varying group sizes, and $\|\beta_l\|$ denotes the l_2 (Euclidean) norm of the coefficients, which is not squared. Thus, instead of constraining the optimisation by the sum of the absolute value of individual coefficients, group lasso constrains by the l_2 norm of groups of coefficients. Like in lasso, depending on the value of λ, entire groups of predictor coefficients may be set to zero.

A second challenge with lasso is that variable selection can be inconsistent and that many noise variables can be included in the estimate, especially with increasingly large p. Meinshausen and Bühlmann (2006) and Zhao and Yu (2006) show that this shortcoming leads to conflict between optimal prediction and consistent variable selection. They show that the optimal λ for prediction can give inconsistent variable selection results, including noise variables in the model and biased estimates for large coefficients. These studies confirm that under certain conditions, lasso does not possess the 'oracle' property. In the context of linear regression, a method possesses the oracle property if it consistently and correctly selects the non-zero coefficients, and their estimates are the same as they would be if the zero coefficients were known in advance (Fan and Li, 2001).

Zou (2006) confirm that there are scenarios in which lasso selection cannot be consistent and thus does not possess the oracle property. They propose a new version of the lasso, the adaptive lasso (*adLasso*), which addresses this problem. It does so by including adaptive weights to penalise different coefficients in the l_1 penalty. Estimates for adaptive lasso are given by:

$$\hat{\beta}^{adLasso} = \underset{\beta}{\operatorname{argmin}} \left(RSS(\beta) + \lambda \sum_{j=1}^{p} \hat{w}_j \mid \beta_j \mid \right) \tag{4.6}$$

where $\hat{w}_j$ is a vector of adaptive weights assigned to the different coefficients. The weights vector is defined as:

$$\hat{w}_j = \frac{1}{(|\ \hat{\beta}_j\ |)^{\gamma}} \tag{4.7}$$

where $\hat{\beta}_j$ here is an initial estimate of coefficients, usually either $\hat{\beta}^{OLS}$, $\hat{\beta}^{ridge}$, or $\hat{\beta}^{lasso}$, and γ is a positive constant for adjustment of the adaptive weights vector. Zou (2006) suggest values of 0.5, 1, and 2. Given this additional weights multiplier, adaptive lasso penalises coefficients with lower initial estimates more than it does larger coefficients. The authors explain that in $p \gg n$ situations, l_2 regularisation can be used to compute the initial estimates of coefficients, given that both l_1 regularisation and OLS are not appropriate estimators in high-dimensional settings. This paper uses $\hat{\beta}^{OLS}$ estimates and $\gamma = 1$ for the adaptive weights vector.

4.3.3 Model training, selection, and validation

All five of these regularisation methods are applied to multivariate survey and annual electricity consumption data for a large sample of U.S. households in California. In addition, a stepwise regression method is also applied for comparison purposes. This section explains the procedures to train the models, select models using cross-validation, and test these on hold-out data.

Models for each regularisation method are trained on a sample of 80% of the observations, subsequently referred to as the 'training set', with 20% held out as the 'test set'. Motivations for this split and for holding out the test set are given in Friedman et al. (2001).

Whereas best-subset and stepwise methods use test statistics for model selection, these are less appropriate for regularisation methods. Instead, cross-validation is used for model selection. In cross-validation, we fit the model using a sample of 90% of the observations and then use it to predict with the remaining 10% of the data in order to obtain the mean-squared error (MSE). This is repeated for k (usually 10) 'folds', and the MSE is averaged over these folds.

In regularisation, it is typical to use cross-validation as a means of computing model

MSE for a range of different λ values in order to see how increases in the strength of the penalty term relate to trade-offs between the bias and variance of the model. Plotting the relationship between λ and MSE obtained through 10-fold cross-validation enables the modeller to select a final model that minimises MSE or (given the objective of the analysis) select a more parsimonious model that still gives an MSE within one standard error of the minimum. This last piece of guidance is given in Friedman et al. (2010, p. 17). In elastic net regularisation, both λ and α are tuned simultaneously across a range of values to find the combination of l_1 and l_2 penalties that minimises the MSE.

After selecting a model for a given value of λ, the model is applied to the test set, and fitted values for the response are compared with actual values to determine prediction error. To evaluate each of the regularisation models, we compare three criteria: model root mean-squared error (RMSE), R-squared for the test set, and the number of non-zero coefficients. These criteria permit an evaluation of the competing aims of prediction and sparsity for the models.

There are several reasons why typical inferential constructs such as confidence intervals and p values are not calculated in this analysis. One reason is that inference is not entirely appropriate given the non-random sample of households studied. Additionally, these inferential constructs do not generally exist for penalised regression estimates. Taylor and Tibshirani (2015) state this problem in simple terms: if we use a regularisation method for variable selection, we have already searched for the strongest associations in the data and selected these. This means the bar for declaring associations significant must be set higher. There is an emerging literature on post-selection inference (Lee et al., 2016; Lockhart et al., 2014; Taylor and Tibshirani, 2015), but in this paper, given the nature of the sample and the aim to identify and describe factors that have strong associations with electricity usage, inference is not part of the analysis.

4.3.4 Multiple imputation for missing data

The penalised regression approaches introduced in the previous section can improve the performance of statistical models when many of the challenges discussed are present. The final challenge mentioned in the introduction was that of missing data, which is not addressed through regularisation.

Issues of missing data are well documented and quite common in social science research. Several authors conducted a review of literature employing surveys in political science journals and found that "approximately 94% use listwise deletion to eliminate entire observations (losing about one-third of their data, on average) when any one variable remains missing…" (King et al., 2001, p. 45). Statisticians and methodologists agree that this is a poor approach to handling missing data because it can both result in the loss of information and introduce bias into regression models (Pampaka et al., 2016). In the case of this study, 126 full observations would have been deleted following this approach (a loss of 13% of the data). For this reason, a multiple imputation (MI) method is used to handle missing data.

Multiple imputation (MI) extracts information from the observed variables with a statistical model (for instance, a linear model), uses the model to predict multiple values for each missing data point, and then uses these to construct multiple complete data sets (Honaker and King, 2010; Rubin, 1996). In each imputed data set, the observed values are the same while the imputed values vary based on the uncertainty in predicting each missing value. The analysis can then proceed as it normally would on each of these full data sets, afterwards combining or 'pooling' the results.

Improved computational power has made MI relatively easy to implement. This paper uses the Expectation-Maximisation with Bootstrapping (EMB) method to create and implement an imputation model with m data sets. For the sake of brevity, algorithmic details are not included here, but they are available in detail in Honaker et al. (2011) and Takahashi (2017).

Missing values are assumed to be missing at random (MAR), meaning "the probability of missing data on a particular variable may depend on other observed variables (but not itself)" (Pampaka et al., 2016, p. 22). This differs from missing completely at random (MCAR), where missing data are missing due to random error, and not missing at random (NMAR), where missing data are due to respondents refusing to answer questions for specific reasons, and these answers cannot be predicted from the other data. A relevant example is household income, where refusal to answer this question may be non-random. This analysis assumes income can reliably be predicted from other variables in the data, such as age, size of dwelling, and others.

4.3.5 Software

The statistical software *R Statistics* is used for all analyses (R Core Team, 2017). Ridge, lasso, and adaptive lasso are all computed using the *glmnet* package (Taylor and Tibshirani, 2015). Group lasso is computed using *gglasso* (Yang and Zou, 2017), and elastic net is computed using *caret* (Kuhn, 2008). This is also the package used to evaluate final models. *MASS* is used to compute a stepwise regression model for comparison purposes (Ripley and Venables, 2002). Finally, *Amelia II* handles multiple imputation (Honaker et al., 2011).

4.4 Data

The data for this study come from a detailed survey and a database of electricity usage for utility customers in Palo Alto, California. This section introduces the data collection procedures and then presents tables of descriptive statistics for all variables.

4.4.1 Palo Alto residential profile

Palo Alto is a city in the California Bay Area. It has 66,500 residents, a mild, Mediterranean climate, and a five-year (2014–2018) average of 1,216/413 heating/cooling degree days (base 18°C) (WRCC, 2006). The city has a target of reducing emissions 80% by 2030 and has already achieved reductions of 36% from 1990 levels (City of Palo Alto, 2016). It aims to achieve 16% of these reductions from reducing energy use in existing homes.

Palo Alto's average residential electricity use is 529 kWh per month, similar to the state-wide average of 557 kWh (CEC, 2017) but well below the U.S. average of 900 kWh (U.S. EIA, 2016).

4.4.2 Data collection

A detailed household survey was delivered by e-mail to customers of the city's municipal utility, City of Palo Alto Utilities, which is the sole provider of electric, gas, and water utilities for most of the city's residents. The survey covers 56 questions on occupant socio-demographics, physical dwelling characteristics, occupant attitudes toward the environment, knowledge of energy issues, occupant curtailment behaviours, and energy efficiency programme participation. Survey questions were refined with the help of a focus group of the utility's customers.

Utility customers for whom an e-mail address was on record received an invitation to participate. Of 11,963 emails, 4,639 were opened and 1,247 surveys were completed. The completion rate was 8% without incentive, 15% when offered entry into a lighting retrofit lottery, and 27% when offered an LED light bulb.

Historical billing data for all of the utility's customers was shared with the researcher. Households were excluded from the analysis if their home address was incomplete or did not match the utility records (79 cases). Households with PV installations were removed to avoid erroneous use of net-demand data (160 cases), leaving $N = 1008$ for use in this analysis.

4.4.3 Independent variables

Independent variables are grouped into categories matching those reviewed in the literature. In the tables below, variables are presented along with their coded numerical ranges and descriptive statistics (where variables are continuous, means are presented as 'M' and standard deviations as 'SD'; for categorical variables, the categories in bold indicate reference categories for regression analyses).

Tables 4.1 and 4.2 show the socio-demographic variables and characteristics of dwellings in the sample. Where data are available, these tables also include variable frequencies for the full city-wide population from the American Community Survey (ACS) (U.S. Census Bureau, 2015). Overall, the sample is a good representation of the Palo Alto population, while property owners, elderly households, and detached dwellings are overrepresented.

Description (codes)	**Response**	**Sample frequency**	**Population frequency**
Gender (0–1)	**Female**	39%	49%
	Male	61%	51%
Age range (1–8)	18–25	<1%	29%
	26–35	3%	12%
	36–45	10%	14%
	46–55	22%	15%
	56–65	25%	11%
	66–75	25%	9%
	76–85	12%	5%
	86 and above	2%	3%
Highest level of education obtained (1–3)	Some college or less	4%	20%
	College graduate (four-year degree)	26%	28%
	Postgraduate	69%	52%
Tenure: own or rent (1–2)	**Own**	87%	55%
	Rent	13%	45%
Number of occupants (1–5)		$M = 2.53$, $SD = 1.13$	$M = 2.53$
Total household income before taxes during past 12 months (1–5)	<\$50,000	8%	20%
	\$50,000–\$99,999	16%	18%
	\$100,000–\$199,999	34%	28%
	\$200,000–\$499,999	34%	34%†
	>\$500,000	8%	
Electric rate schedule (1–2)	**Regular electric**	98%	
	Time-of-use	2%	

Note: Population $N = 66478$. Population data are from the American Community Survey (ACS) (U.S. Census Bureau, 2015).
† Includes '\$200,000 and above'.

Table 4.1: Summary and descriptive statistics for socio-demographic variables.

Energy literacy is assessed with the questions in Table 4.3. Correct responses to each question are bolded in the table. These items are based on similar work by Coyle (2005), Brounen et al. (2013), DeWaters and Powers (2013), DeWaters et al. (2013), and Southwell et al. (2012). On average, participants scored 4 out of 7 possible correct answers.

Table 4.4 shows the occupant attitude variables and frequencies. These are either

Description (codes)	**Response**	**Sample frequency**	**Population frequency**
Size range of home in square feet (1–6)	Less than 1000	11%	
	1001-1500	23%	
	1501-2000	29%	
	2001-2500	19%	
	2501-3000	10%	
	More than 3000	8%	
Year of construction (1–3)	Pre-1950	28%	23%
	1950-1989	57%	59%
	1990–present	15%	18%
Type of home (1–2)	**Attached or apartment building**	20%	56%
	Detached home	80%	44%
Number of bedrooms (1–5)	1 or 2	19%	42%
	3	36%	31%
	4	33%	20%
	5 or more	11%	7%
Home has double- or triple-glazed windows (0–1)†	**No**	30%	
	Yes	70%	
Home has floor insulation (0–2)	Not sure	21%	
	No	54%	
	Yes	25%	
Home has roof insulation (0–2)	Not sure	10%	
	No	14%	
	Yes	76%	
Home has wall insulation (0–2)	Not sure	18%	
	No	24%	
	Yes	58%	
Presence or not of an air conditioning system (0–1)	**Does not have AC**	61%	
	Has AC	39%	
Energy devices present in the home (0–1)	Solar water heating	4%	
	LED lighting	77%	
	Smart meter	5%	
	Wi-Fi thermostat	14%	
	Programmable thermostat	58%	
	Plug-in electric vehicle	13%	
	In-home energy display	2%	
	Other energy device	2%	

Note: Population $N = 27555$ households. Palo Alto data are from the ACS (U.S. Census Bureau, 2015).
† 'Not Sure' combined with 'No' responses given low frequencies in these categories.

Table 4.2: Summary and descriptive statistics for physical dwelling variables.

measured on a 5-point Likert scale ('Strongly disagree' to 'Strongly agree') or are dummy coded. The final variable in this section is binary coded and thus measures whether respondents believe renewable energy is beneficial primarily for environmental impact (1) or other reasons (0). The mean correlation coefficient between all attitude variables is $r = 0.17$. The Likert scale questions show slightly stronger correlations, with a mean correlation coefficient of $r = 0.29$.

Behavioural variables are shown in Table 4.5. Except for the efficiency and rebate

Energy literacy quiz question	**Response option**	**Frequency**
1. Who owns your utility company?	A private entity	0.7%
	The State of California	0.3%
	City of Palo Alto	97%
	Pacific Gas & Electric (PG&E)	2%
2. How much do you pay per kWh for electricity?	Less than 5 cents	6%
	5 - 10 cents	25%
	11 - 20 cents	56%
	21 - 30 cents	8%
	More than 30 cents	5%
3. How much electricity do you think an average Palo Alto single-family household consumes each month?	0 to 10 kilowatt-hours (kWh)	1%
	11-100 kWh	9%
	101-500 kWh	42%
	501-1,000 kWh	41%
	1,001-5,000 kWh	7%
4. Which of the following resources generates the most electricity in California?	Oil	8%
	Coal	5%
	Natural gas	49%
	Nuclear	4%
	Hydroelectric	29%
	Solar	3%
	Wind	2%
5. What percentage of the electricity supplied by City of Palo Alto Utilities is carbon neutral?	20	16%
	30	23%
	50	19%
	70	16%
	100	25%
6. Which of the following uses the most energy in the average Palo Alto home over the course of a year?	Lighting	4%
	Powering household appliances	14%
	Heating water	11%
	Heating and cooling rooms	63%
	Refrigerating food	9%
7. Of the following household appliances, which do you think consumes the most electricity while being used?	Dishwasher	7%
	Fridge/freezer	19%
	Laptop computer	2%
	LED light bulb	1%
	Electric space heater	71%

Table 4.3: Energy literacy quiz questions and response frequencies.

Description (codes)	**Mean (SD) or Response (frequency)**
Saving energy is important	4.62 (0.64)
I would do more to save if I knew how	3.81 (0.88)
We don't have to worry about conserving energy because new technologies will be developed to solve problems†	4.17 (0.91)
California should produce more electricity from renewables	4.39 (0.80)
Laws protecting the natural environment should be made less strict to produce more energy†	4.03 (1.08)
The way I personally use energy does not really make a difference to the energy problems in California†	3.74 (1.02)
My decisions to participate in energy efficiency programmes are mostly driven by the amount of money I can save†	2.91 (1.10)
Renewable energy is still too expensive to be practical for California†	3.55 (1.09)
When you think about energy, what are the most important values to you? (0–1)	Comfort (46%) Ease of use (31%) Expense (71%) Safety and security (49%) Ability to go off-grid (7%) Environmental stewardship and protection (67%)
What do you see as the most important benefit of renewable energy? (0–1)	**Reducing impact on environment (80%)** Reducing personal energy costs (7%) Decreasing dependence on foreign energy imports (6%) Helping support 'green' job creation (2%) Enabling off-grid capabilities (2%) I do not see any benefits to renewable energy (1%) Other (2%)

† Likert scale is reverse coded.

Table 4.4: Summary and descriptive statistics for occupant attitude variables.

variables, which are measured as continuous predictors, these variables are measured on a 3-point Likert scale ('Never', 'Sometimes', 'Always'). Correlations are generally low, with a mean correlation coefficient of $r = 0.11$. While the means for the curtailment variables indicate high frequencies of energy saving behaviour, especially curtailing AC use, both energy efficiency programme participation and rebate uptake are low.

4.4.4 Missing data

Table 4.6 shows the frequency of missing data for the socio-demographic variables, which were made optional on the survey. Household income has the most missingness, while missingness amongst the other data is generally low.

After specifying these variables and setting their logical bounds from the variable codes,

Description (codes)	**Mean (SD)**
How often do you...	
Turn off lights and electrical appliances when not in use	1.70 (0.47)
Unplug electrical appliances when not in use for an extended period	0.87 (0.71)
Take a shorter shower to conserve energy used for heating water	1.33 (0.66)
Purchase appliances that are ENERGY STAR® or energy efficiency labeled	1.58 (0.56)
Only run the dishwasher or clothes washer/dryer when full	1.75 (0.50)
Turn down thermostat while asleep in the winter†	1.76 (0.54)
Turn off AC when no one is home in the summer†	1.86 (0.37)
Talk with other members of your household about your energy bill†	1.49 (0.72)
Talk with your friends or neighbours about your energy bill	0.78 (0.76)
Talk with your friends or neighbours about ways to conserve energy	1.02 (0.76)
Talk with your friends or neighbours about your own energy efficient devices or technologies	1.01 (0.78)
Number of energy savings programmes respondent participated in (0–3)‡	0.63 (0.82)
Number of energy rebates respondent has received (0–3)‡	0.40 (0.74)

† N/A response frequencies (*TurnDownTherm* = 29; *TurnOffAC* = 618; *TalkAboutBillFam* = 79).

‡ Includes '3 and above'.

Table 4.5: Summary and descriptive statistics for occupant behaviour variables.

Variable	Missing (%)
Gender	2
Age	2
Education	1
Tenure	1
Occupancy	1
Income	13

Table 4.6: Missing value frequencies for socio-demographic variables ($N = 1008$).

MI using the EMB method described in Section 4.3.4 was used to impute five completed data sets.[2] Plots for each of the socio-demographic variables across these five sets were inspected and compared with the original data, which show similar distributions, thus providing a degree of validation. Distributions for the *Income* variable across the five imputed data sets can be found in the appendix (Figure 4.A.1).

Again, because listwise deletion would reduce the number of observations by 13%, the goal for imputation is to avoid losing this important information when conducting subsequent analyses. The total missingness of the data is low, however, so the additional step of 'pooling' results across the five imputed sets is not taken due to computational complexities. Instead, one of the five imputed data sets is randomly selected and used in all subsequent analyses.

4.4.5 Dependent variable: Annual electricity consumption

The dependent variable for the regression analyses is 2016 annualised electricity consumption in kilowatt-hours (kWh). Table 4.7 shows electricity usage summary statistics for both the sample and the utility's full customer population. The sample includes 12 households with more than 18,000 kWh for the year. The correct operation of their meters was validated, and they are kept in the sample.

	N	Mean	SD	1st Quantile	Median	3rd Quantile
Sample	1,008	6,116	3,656	3,759	5,449	7,585
Population	20,006	6,040	5,596	3,130	4,930	7,430

Table 4.7: Electricity usage summary statistics for sample and customer population.

To address the heteroscedasticity of regression errors, the dependent variable is log-transformed prior to analysis. While the log-transformed electricity usage distribution still exhibits some skew, it is more normally distributed. The log transformation is chosen to improve the regression residuals while still enabling a relatively simple interpretation of

[2]The authors of *Amelia II* recommend a standard value of $m = 5$ (Honaker et al., 2011).

results.[3] The sample's mean electricity consumption before transformation is $M = 6166$ with a standard deviation of $SD = 3656$. Considering the wider California Bay Area, the mean annual electricity consumption across eight Bay Area counties in 2015 was 6,096 kWh (CEC, 2017; U.S. Census Bureau, 2015).

4.5 Results

The five regularisation methods introduced in Sections 4.3.1—4.3.2 are applied to the data set of household survey responses and log-transformed annual electricity consumption. A stepwise regression is also computed to provide some comparison between regularisation and other variable selection techniques. The total number of predictors included in the data is 58.

Figures 4.2–4.3 show the results for 10-fold cross-validation to tune the penalty parameter λ and select an optimal model using each regularisation method. These plots show how cross-validation MSE varies as a function of the penalty parameter. High bias models are expected on the right side of these plots where the values of λ are higher, whereas high variance models are expected on the left side where the values of λ are lower. In the cases shown here, the characteristic U-shape of the the bias-variance trade-off is very slight (and in some cases absent altogether). This suggests the models are not overfitting much, even with very small amounts of regularisation. This may be due to the number of predictors being large but not in comparison to the number of observations. The plots do, however, show that models with large penalties have high bias and greater cross-validation MSE as a result.

The plots also show that there is not a sizeable difference in the regularisation paths for the five methods, and each is able to achieve similar minimum cross-validation MSE

[3]The dependent variable changes by 100 $\times$ (coefficient) percent on average for each one unit increase in the predictor variable while all other predictor variables are held constant. If the predictor is a dummy variable, when its value switches from 0 to 1, the percent change of the dependent variable is $[100(e^{B_1} - 1]$ while the reverse is $[100(e^{-B_1} - 1]$ (Halvorsen and Palmquist, 1980).

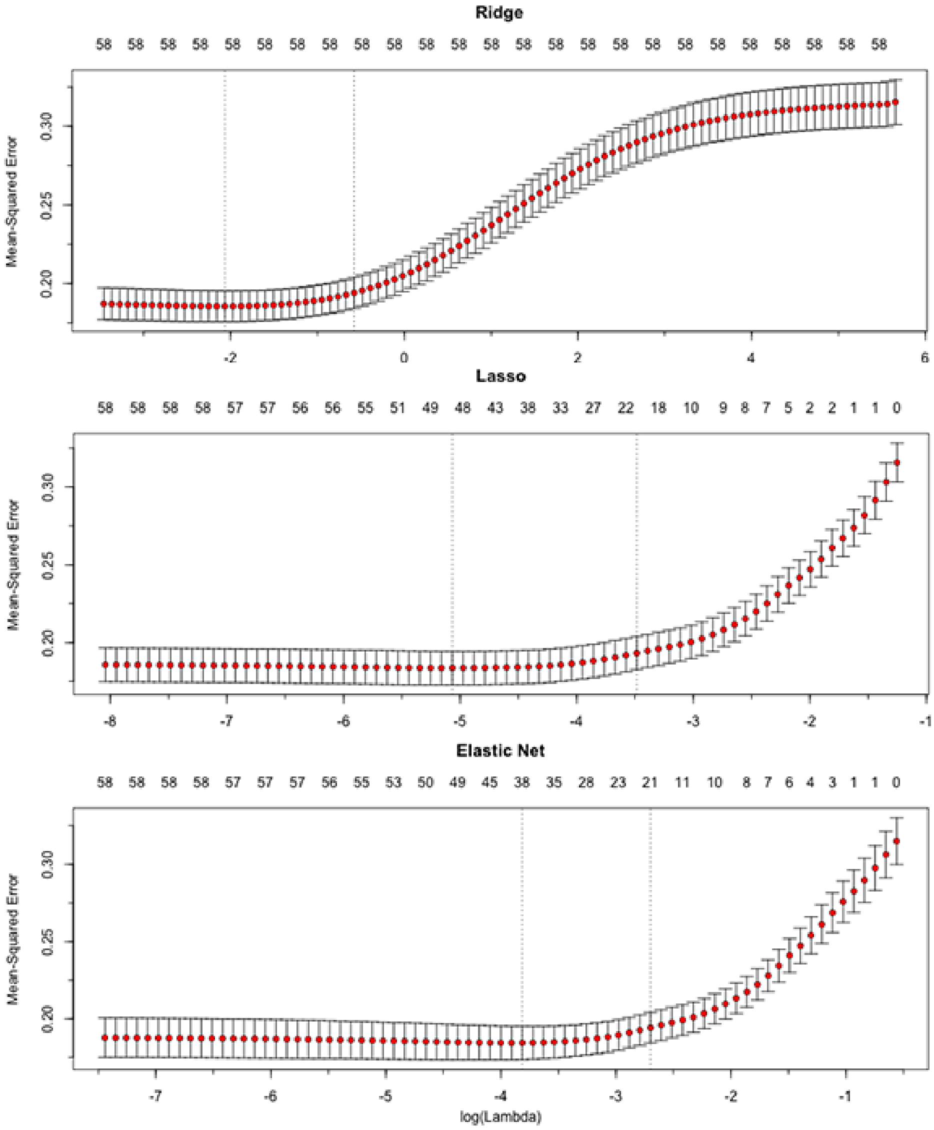

Figure 4.2: Plots of cross-validation MSE for ridge, lasso, and elastic net models. The horizontal bottom axis shows the logarithm of the tuning parameter λ, while the top horizontal axis shows the number of non-zero coefficients in each model. Points and error bars represent the mean and standard error of cross-validation MSE, respectively. The vertical dotted lines give the model with the minimum MSE (left) and with the fewest non-zero coefficients within one standard deviation of the minimum MSE (right).

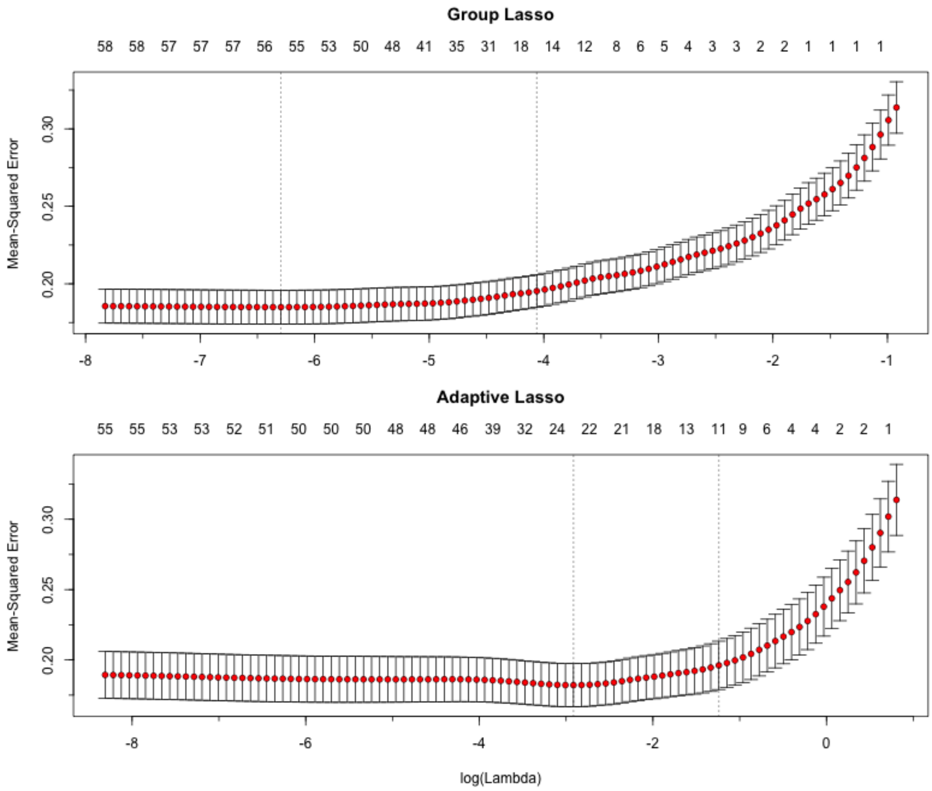

Figure 4.3: Plots of cross-validation MSE for group and adaptive lasso. Axes and plot elements are the same as in Figure 4.2.

(albeit at different strengths of the penalty parameter).

The main difference between the methods, then, can be seen in their sparsity or number of non-zero coefficients, which is indicated along the top horizontal axis. While ridge regression does not set any variable coefficients to zero, retaining all 58 predictors in the final model, both lasso and elastic net achieve similar levels of sparsity, though elastic net reaches a model with minimum MSE needing 20% fewer predictors than lasso (left vertical dotted lines). For the most sparse models that have a cross-validation MSE within one standard error of the minimum (right vertical dotted lines), elastic net and lasso methods both select 21 predictors, which is a 60% reduction from the original total.

Group and adaptive lasso select even more parsimonious models. The plots in Figure 4.3 show that group lasso finds a model within one standard error of the minimum containing

16 predictors, while adaptive lasso selects a model containing just 11 predictors, a reduction of 81% of the original predictor set.

In order to compare the performance of these methods with other variable selection approaches, a forward stepwise regression is computed using AIC as the criteria for model selection. Next, each of these six models is applied to the test set. For all regularisation models, the model within one standard error of the minimum is the model used on the test data. The rationale for this is that selecting the most parsimonious model across each method permits comparisons between predictive error and model interpretability, which is the key objective of this analysis.

Models are compared across several criteria, including RMSE and R-squared for predictions given the test data, as well as the number of non-zero coefficients in the model. RMSE is measured in units of the dependent variable, in this case log-transformed annual electricity consumption, which has a mean of 8.57 and a standard deviation of 0.56. Table 4.8 shows results for all six methods, ordered by increasing RMSE. The table shows that elastic net and group lasso achieve the lowest test data RMSE. The stepwise method gives a lower test set error than either adaptive lasso or lasso, while ridge gives the highest error. In general, however, the errors are narrowly distributed, signalling not much difference between methods for prediction (which is similar to the results from cross-validation). Note that R-squared is not adjusted, meaning this value does not take into account the number of predictors in the model. Those models with more coefficients would have lower adjusted R-squared values.

Variables selected by elastic net, group lasso, and adaptive lasso and their standardised coefficients are shown in Table 4.9. Both elastic net and lasso select the same predictors without much variation in their coefficients, while both group lasso and adaptive lasso achieve more parsimonious models, which is why variables in these models are inspected. The table separates predictors of high and low usage and lists these in order of standardised coefficient magnitude (averaged across the three model selection methods). Standardised coefficients are measured in the original units of each independent variable and thus cannot

Method	RMSE	R-squared	Non-zero Coefficients
Elastic Net	0.4457	0.3959	21
Group Lasso	0.4475	0.3893	16
Stepwise	0.4481	0.3978	27
Adaptive Lasso	0.4535	0.3789	11
Lasso	0.4538	0.3733	21
Ridge	0.4562	0.3751	58

Table 4.8: Model RMSE, R-squared, and number of non-zero coefficients for methods applied to test data.

be accurately compared with coefficients of other independent variables measured on different scales. Because most of the variables included in this model are measured in different units (e.g., *Age* in years and *Income* in dollars), standardised coefficients allow for a comparison of each predictor's relative importance in explaining household electricity consumption.[4]

The positive predictors selected by all three methods are EV ownership, size of home, home occupancy levels, type of home, and AC ownership. Other positive predictors include valuing comfort in relation to energy use, household income, solar water heating, being on a time-of-use rate, and respondent age.

The frequency with which respondents report turning off AC when not home, unplugging appliances when not in use, and taking a shorter shower to save energy are associated with lower use, as is renting versus owning a home. Because less than half of the study sample reported ownership of AC, to examine the effect of the AC curtailment variable, the three methods are used to fit models to survey data for only those households that own AC ($N = 390$). Both elastic net and adaptive lasso select models including the variable measuring AC curtailment. Cross-validation curves for these models are included in the appendix (Figure 4.A.2).

[4]Standardised regression coefficients are measured in standard deviations rather than in the original units of the independent variable, so the coefficient indicates the number of standard deviation changes expected in the dependent variable for a one standard deviation change in the independent variable.

Other behavioural and attitudinal variables exhibiting a negative relationship with consumption have relatively small standardised coefficients.

Of these selected variables, nine measure characteristics of the dwelling and appliance or energy-related device ownership. Seven measure occupant socio-demographics, six measure attitudinal factors, and six measure behavioural factors. For the top ten variables with the strongest associations to electricity consumption, seven are either dwelling characteristics or socio-demographics, two are behavioural variables, and one is an attitude variable.

Predictor	$\hat{\beta}_{elastic}$	$\hat{\beta}_{gLasso}$	$\hat{\beta}_{adLasso}$
High usage predictors			
EV ownership	0.155	0.102	0.214
Size of home	0.115	0.136	0.138
Occupancy level	0.080	0.109	0.089
Type of dwelling	0.098	0.036	0.109
Important value: comfort	0.053	0	0.055
Household income	0.039	0.063	0
Solar water heating	0.028	0	0.064
Time-of-use rate	0.023	0	0.057
AC ownership	0.025	0.021	0.018
Age	0.002	0.041	0
Roof insulation	0.013	0.027	0
Number of bedrooms	0.023	0.016	0
Other device ownership	0	0	0.023
Renewables too expensive	0	0.015	0
Talk with family about bill	0.001	0.011	0
Smart thermostat ownership	0.011	0	0
Gender	0.010	0	0
Talk with family about conservation	0	0.003	0
Low usage predictors			
Behaviour: turn off AC	−0.094	0	−0.186
Behaviour: unplug appliances	−0.075	−0.089	−0.062
Rents home	−0.071	0	−0.077
Behaviour: take a shorter shower	−0.006	−0.023	0
Important value: cost of energy	−0.019	0	0
New technologies will solve problems	0	−0.018	0
Benefit of renewables	−0.010	0	0
Behaviour: turn off lights	−0.008	0	0
Would do more to save if I knew how	0	−0.001	0

Table 4.9: Variables selected across elastic net, group lasso, and adaptive lasso models. Variables are split by the sign of their effect and are ordered by the magnitude of their standardised coefficient.

4.6 Discussion

4.6.1 Summary of results and comparison to previous research

The methodological results of this study show that the regularisation methods introduced in this paper achieve an RMSE on the test set ranging from 0.4562–0.4457, which is equivalent to 0.81–0.79 of the response variable's standard deviation. In other words, the prediction error of these methods is 20% smaller than the standard deviation of log-transformed annual electricity consumption. These prediction error results compare favourably with those of other studies employing regularisation methods for building energy consumption prediction (Deng et al., 2018; Hsu, 2015; Robinson et al., 2017). Across methods, the other goodness-of-fit measure, R-squared (ranging from 0.375–0.396) is consistent with or surpasses results of many previous studies of household electricity usage (Abrahamse and Steg, 2009; Brandon and Lewis, 1999; Huebner et al., 2016; McLoughlin et al., 2012; Wahlström and Hårsman, 2015; Wiesmann et al., 2011).

Returning to the question of comparing model predictive accuracy with sparsity, the regularisation methods (excluding ridge regression) yield a sizeable reduction in the number of variables needed to achieve similar predictive accuracy. Comparing adaptive lasso and stepwise regression, for instance, adaptive lasso selects a model that has less than a 2% greater prediction error than stepwise regression but reduces the number of variables in the model by a further 27%. Trading off a small increase in prediction error for a large reduction in the number of variables that needs to be collected is favourable when the objective of analysis includes a simpler, more interpretable model.

The empirical results of this study confirm that the size of home and number of occupants are two of the strongest determinants of residential electricity use patterns (Brounen et al., 2012; Huebner et al., 2016; Jones et al., 2015; Kavousian et al., 2013; Wyatt, 2013).

Type of dwelling (Huebner et al., 2016; Wyatt, 2013; Yohanis et al., 2008) and household income (Cramer et al., 1985; Brounen et al., 2012; Yohanis et al., 2008) can be confirmed as strong predictors even though results from other studies are mixed on their effect (e.g., Bedir et al., 2013; Kavousian et al., 2013). Previous findings on the associations between tenure type and electricity consumption are similarly mixed, with some supporting this study's findings of higher consumption in privately-owned residences (Wiesmann et al., 2011; Wyatt, 2013; Yohanis et al., 2008), while others report either higher consumption in rented buildings or no significant effect (Bedir et al., 2013; Kavousian et al., 2013; Ndiaye and Gabriel, 2011; Tso and Yau, 2007).

Unsurprisingly, EV ownership and presence of AC in the home both exhibit positive associations with annual electricity use. Given the growing uptake of these technologies around the world, a detailed understanding of their impact on total consumption is increasingly important.

Occupant attitudes toward energy conservation and renewable energy, the values occupants consider important in relation to energy use, and their knowledge of energy concepts do not exhibit strong associations with electricity consumption. These results support those of previous studies that do not find a notable link between environmental attitudes and electricity consumption (Bartiaux and Gram-Hanssen, 2005; Brounen et al., 2013; Vringer et al., 2007). Furthermore, rates of energy knowledge as demonstrated by performance on an energy quiz bear little association to electricity usage, which is similar to the findings of Brounen et al. (2013). One attitude variable is particularly strong in comparison to other predictors: listing 'comfort' as one of the most important values related to energy use is associated with higher consumption. This supports Wilhite et al.'s (2000) argument that notions of comfort and convenience may have considerable implications for electricity demand and are not sufficiently addressed in energy demand research.

From the set of behaviour variables, unplugging appliances when not in use for extended periods and turning off AC when not needed are selected as predictors of lower usage in the model. These results confirm those of Wallis et al. (2016), who find a statistically significant association between habitual energy saving behaviours and reduced annual consumption.

Participation in energy efficiency programs and uptake of rebates for efficient appliances are not among the significant predictors in the model. This may be due to very low rates of participation and uptake reported amongst the survey sample.

4.6.2 Implications of results

These results have implications both for statistical approaches for modelling building electricity consumption as well as for understanding factors that influence consumption.

Given the complexities of residential electricity consumption, statistical methods that reduce large predictor sets without sacrificing much predictive accuracy are advantageous in studies of domestic electricity demand. The regularisation methods introduced in this paper, including extensions to the lasso that take into consideration some of its methodological weaknesses, are useful in this regard. Furthermore, these methods are computationally efficient and can address several important statistical challenges, such as model overfitting, multicollinearity, and high-dimensional data. Even in the absence of these issues, the methods presented here can effectively identify key variables in models of building energy consumption, and they do not suffer from the same statistical weaknesses as do other variable selection approaches. For these reasons, they are especially suitable for building energy modelling, give the specific challenges faced in this discipline.

This paper has stressed the importance of letting the analysis objectives and the characteristics of the data guide the use of regularisation methods. It has explained why, for instance, both elastic net and lasso are likely to show similar results given the absence of strong multicollinearity effects and high-dimensionality, and it has confirmed this empirically (lasso and elastic net select the same variables, although prediction error is somewhat higher for lasso). The extensions to the lasso are introduced to improve upon these results, and we see that they do (in terms of yielding simpler models without much loss in predictive accuracy).

The empirical implications of this study are best understood in the context of the study location. Palo Alto's population is projected to grow at a rate of 1.1% annually over the next 20 years. The city's senior population (65 and over) is one of its fastest growing demographics (ABAG, 2015). In this region, large, detached homes are commonplace, occupancy levels are growing, and the city's average median family income is the third highest in the U.S. (Pisillo, 2012). Given the demonstrated effects of dwelling size and type, occupancy levels, and household income on residential electricity consumption, these trends are important to consider when determining ways to meet the city's ambitious energy savings targets.

This study provides evidence that policies or programs that further improve the thermal performance and efficiency of residential buildings are necessary to achieve substantial emissions reductions. This evidence may be especially relevant for single-family, detached homes in Palo Alto. Given the study's findings that energy efficiency programme uptake is low, more effort is needed to engage residential customers in this regard. Encouraging regular home energy audits through building codes and regulations could help determine where home efficiencies are lacking. A target audience for these initiatives should be the city's older residential population, as a link was found between age and electricity consumption in the models.

Despite Palo Alto's relatively mild climate (less than half the sample owned AC), the significance of AC for consumption suggests that reducing AC use deserves special attention. This is even more pressing given the anticipated rise in home AC ownership in middle-income countries with additional warming. Davis and Gertler (2015) predict near-universal saturation of AC in all warm areas in just a few decades, and their findings suggest AC impacts on energy usage will be larger than previously believed. AC adoption and use must be met with even greater energy efficiency gains or behavioural changes to reduce its projected impact, especially considering the effect of AC use on peak demand.

The same applies to EV ownership. Palo Alto has one of the highest rates of EV ownership in the country (around 3–4% of registered vehicles) and aims for 90% of registered vehicles to be electric by 2030. This study provides further evidence that EV ownership must be met with vehicle-to-grid integration projects and smart charging policies to lessen the substantial burden this transformation will place on local electricity networks (Sidhu et al., 2009).

This study provides evidence that households can decrease their electricity usage by engaging more frequently in energy saving behaviours, especially those related to appliances and AC. While 70% of respondents report they 'Always' turn off lights and appliances when not in use, only 20% report the same for unplugging their appliances. Further savings could be achieved given that standby power consumption is responsible for around 15% of household electricity usage in California (LBNL, 2004). Much of the focus in reducing residential electricity consumption has been on deploying energy efficiency measures rather than motivating changes in behaviour, but this study highlights the important role of habitual actions taken to save energy in the home and reaffirms previous findings that these can contribute towards reducing carbon emissions (Dietz et al., 2009).

4.6.3 Limitations

This study has limitations to its design, methods, and data. In terms of its design, the sampling methodology is non-random, and participation is limited to utility customers with emails on record. Some of the biases, such as underrepresentation of renters and people in the 18–35 age group, have been discussed. This means that the findings are not necessarily generalisable. The self-reporting of behaviours and attitudes could mean social desirability bias is present and may have influenced results (Fisher, 1993).

The regularisation methods applied in this study show promise for improving building energy prediction and selecting sparse models that highlight key variables. Their application in this study, however, does not showcase their suitability for addressing other issues, such as multicollinearity and overfitting, since these challenges are muted in the data. The cross-validation results suggest the models do not exhibit high variance, even without applying much regularisation. This is likely because the sample size is large compared to the number of predictors. With a smaller sample size, or with an increasingly large number of predictors, the regularisation methods introduced here are likely to improve in performance, especially in their prediction error on the test set. Some evidence of this is seen when fitting the models to these data including only those households with AC

($N = 390$). Cross-validation curves for these data are slightly more U-shaped (see Figure 4.A.2 in the appendix). One further methodological limitation is that the analysis did not consider interactions between predictors. Evidence from previous research suggests these methods and several extensions can handle high numbers of pair-wise interactions, which could enable further insight into the relationships between drivers of building energy consumption (Hsu, 2015).

Regarding the limitations to the data, additional details on appliance ownership and use may increase the explanatory power of the models and yield deeper insights into how occupant behaviour is associated with electricity consumption. Other specific factors not investigated include more detailed efficiency measures taken in the home, data on the type of AC (central versus window unit), pool ownership, and fuel used for space heating. The last of these may be particularly important, given an estimated 25% of Palo Alto households use electricity for heating (U.S. Census Bureau, 2015). Nine respondents indicated ownership of air source heat pumps on the 'Other' device survey question, but a specific question on fuel used for heating could have revealed the influence of electric heating on annual consumption.

Furthermore, this paper is limited in explaining the drivers of specific electricity end uses, such as space heating and cooling, water heating, or appliances, lighting and electronics, which makes comparing results to other study contexts more difficult (Jones et al., 2015). Similarly, the analyses presented here only consider electricity and not natural gas consumption. The modelling techniques presented could be applied to natural gas usage data for further insight on how to reduce residential building emissions from space heating and cooking.

4.7 Conclusions

This paper discusses the use of regularisation methods in linear regression analysis for improving both prediction and interpretation in residential building energy models. It identifies key challenges in energy modelling and explains how regularisation methods can address these. Next, it demonstrates these methods empirically on multivariate survey and household electricity data for a sample of 1,008 households in Palo Alto, California. It tests a wide range of structural and occupant factors across several distinct variable types to determine those exhibiting the strongest associations with annual electricity use.

The results show that regularisation methods can improve upon traditional variable selection approaches, such as stepwise regression, both in terms of prediction error and model interpretability. Elastic net and group lasso make better predictions on hold-out test data than the other methods while reducing the number of non-zero coefficients in the models. Adaptive lasso selects the most sparse model with 11 predictors, a reduction of over 80%, with only a 1–2% higher prediction error than the other methods.

The analysis finds that household electricity use is best explained through a combination of socio-demographic and physical dwelling characteristics. Size of home, occupancy levels, and ownership of an EV and AC are significantly associated with increased electricity usage. While occupants' attitudes toward the environment and their level of energy knowledge do not generally show strong associations with consumption, this paper does find that specific occupant curtailment behaviours, such as unplugging appliances when not in use for extended periods and turning off AC when no one is home, are strong predictors of lower electricity use.

These findings can inform Palo Alto's energy strategy as it embarks on ambitious usage reduction targets over the next several decades. Results are also informative for other cities and regions that want to understand the key variables influencing consumption, or want to use these to better predict future patterns of consumption.

While the evidence presented here does not refute the importance of improving the structural efficiency of the building stock in order to achieve these targets, it also suggests that occupant factors related to curtailment behaviour are drivers of electricity consumption. This insight is particularly important for designing energy policy in places that expect rapid increases in EV and AC ownership. In Palo Alto, these are expected to be near universal by 2050 (City of Palo Alto, 2016). Reducing home size and occupancy levels are more challenging policy changes to implement than are encouraging energy curtailment behaviours. Of course, understanding the most effective ways to do this is of equal importance and is the subject of much ongoing research. Here, especially, is where other disciplinary approaches that examine the socio-technical structures surrounding behaviours or practices may add the most insight.

Situating these results within a growing body of research on the factors that drive household electricity consumption will contribute to future lines of empirical inquiry in this field. The purpose of future research should be to further investigate the links between the factors identified and tested in this paper as well as to explore any number of additional influential factors that influence household electricity consumption. The methods demonstrated in this paper are applicable to a wide variety of energy and building data, and they can be used successfully in contexts where other statistical methods fail (e.g., where the issues of multicollinearity and high-dimensionality are present). Of particular interest for further research is the application of these methods to high-resolution electricity data. Drivers of electricity consumption across months and years may be different than those that influence daily or hourly consumption patterns. Understanding these differences through the use of regularisation in statistical models can inform strategies for reducing demand on both of these time-scales, which is increasingly important for a low-carbon transition.

Acknowledgements

The author would like to thank the editor and several anonymous reviewers for their detailed comments, which helped clarify the structure and aim of the paper.

Thanks also to the City of Palo Alto Utilities for providing usage data and for facilitating deployment of the survey. Thanks especially to Lacey Lutes and Lisa Benatar for their assistance. Thank you to Phil Grünewald, John Fresen, and Pete Walton for their instructive comments and feedback. Any remaining errors in this article are solely the work of the author.

The author gratefully acknowledges financial support for this research from the City of Palo Alto Utilities, the University of Oxford's Environmental Change Institute, and the Sir Peter Elworthy fund. The author declares no competing interests.

4.8 References

ABAG (2015), San Francisco Bay Area state of the region: Economy, population, housing 2015, Technical report, Association of Bay Area Governments, Oakland, CA.

Abrahamse, W. and Steg, L. (2009), 'How do socio-demographic and psychological factors relate to households' direct and indirect energy use and savings?', *Journal of Economic Psychology* **30**(5), 711–720.

Abrahamse, W., Steg, L., Vlek, C. and Rothengatter, T. (2007), 'The effect of tailored information, goal setting, and tailored feedback on household energy use, energy-related behaviors, and behavioral antecedents', *Journal of Environmental Psychology* **27**(4), 265–276.

Amasyali, K. and El-Gohary, N. M. (2018), 'A review of data-driven building energy consumption prediction studies', *Renewable and Sustainable Energy Reviews* **81**, 1192–1205.

Assimakopoulos, V. (1992), 'Residential energy demand modelling in developing regions: The use of multivariate statistical techniques', *Energy Economics* **14**(1), 57–63.

Barr, S., Gilg, A. W. and Ford, N. (2005), 'The household energy gap: Examining the divide between habitual- and purchase-related conservation behaviours', *Energy Policy* **33**(11), 1425–1444.

Bartiaux, F. and Gram-Hanssen, K. (2005), Socio-political factors influencing household electricity consumption: A comparison between Denmark and Belgium, *in* 'Proceedings of the 2005 ECEEE Summer Study on Energy Efficiency in Buildings', European Council for an Energy Efficient Economy, Presqu'île de Giens, France, pp. 1313–1325.

Bedir, M., Hasselaar, E. and Itard, L. (2013), 'Determinants of electricity consumption in Dutch dwellings', *Energy and Buildings* **58**, 194–207.

Belaïd, F., Roubaud, D. and Galariotis, E. (2019), 'Features of residential energy consumption: Evidence from France using an innovative multilevel modelling approach', *Energy Policy* **125**, 277–285.

Brandon, G. and Lewis, A. (1999), 'Reducing household energy consumption: A qualitative and quantitative field study', *Journal of Environmental Psychology* **19**(1), 75–85.

Brohus, H., Heiselberg, P. K., Simonsen, A. and Sørensen, K. C. (2010), Influence of occupants' behaviour on the energy consumption of domestic buildings, *in* 'Clima 2010: 10th Rehva World Congress: Sustainable Energy Use in Buildings', Clima 2010: 10th Rehva World Congress, Antalya, Turkey.

Brounen, D., Kok, N. and Quigley, J. M. (2012), 'Residential energy use and conservation: Economics and demographics', *European Economic Review* **56**(5), 931–945.

Brounen, D., Kok, N. and Quigley, J. M. (2013), 'Energy literacy, awareness, and conservation behavior of residential households', *Energy Economics* **38**, 42–50.

Cayla, J.-M., Maizi, N. and Marchand, C. (2011), 'The role of income in energy consumption behaviour: Evidence from French households data', *Energy Policy* **39**(12), 7874–7883.

CEC (2017), 'Total electric system generation', `http://www.energy.ca.gov/almanac/electricity_data/total_system_power.html`.

City of Palo Alto (2016), Palo Alto Sustainability and Climate Action Plan: Framework, principles, guidelines, goals & strategies, Technical report, City of Palo Alto, Palo Alto, CA.

Coyle, K. (2005), 'Environmental literacy in America: What ten years of NEETF/Roper research and related studies say about environmental literacy in the US', *National Environmental Education & Training Foundation* .

Cramer, J. C., Miller, N., Craig, P., Hackett, B. M., Dietz, T. M., Vine, E. L., Levine, M. D. and Kowalczyk, D. J. (1985), 'Social and engineering determinants and their equity implications in residential electricity use', *Energy* **10**(12), 1283–1291.

Darby, S. (2006), The effectiveness of feedback on energy consumption, Technical report, Environmental Change Institute, University of Oxford.

Davis, L. W. and Gertler, P. J. (2015), 'Contribution of air conditioning adoption to future energy use under global warming', *Proceedings of the National Academy of Sciences* **112**(19), 5962–5967.

Deb, C. and Lee, S. E. (2018), 'Determining key variables influencing energy consumption in office buildings through cluster analysis of pre- and post-retrofit building data', *Energy and Buildings* **159**, 228–245.

Deng, H., Fannon, D. and Eckelman, M. J. (2018), 'Predictive modeling for U.S. commercial building energy use: A comparison of existing statistical and machine learning algorithms using CBECS microdata', *Energy and Buildings* **163**, 34–43.

DeWaters, J. and Powers, S. (2013), 'Establishing measurement criteria for an energy literacy questionnaire', *The Journal of Environmental Education* **44**(1), 38–55.

DeWaters, J., Qaqish, B., Graham, M. and Powers, S. (2013), 'Designing an energy literacy questionnaire for middle and high school youth', *The Journal of Environmental Education* **44**(1), 56–78.

Diamantopoulos, A. (2006), 'The error term in formative measurement models: Interpretation and modeling implications', *Journal of Modelling in Management* **1**(1), 7–17.

Dietz, T., Gardner, G. T., Gilligan, J., Stern, P. C. and Vandenbergh, M. P. (2009), 'Household actions can provide a behavioral wedge to rapidly reduce U.S. carbon emissions', *Proceedings of the National Academy of Sciences* **106**(44), 18452–18456.

Djuric, N. and Novakovic, V. (2012), 'Identifying important variables of energy use in low energy office building by using multivariate analysis', *Energy and Buildings* **45**, 91–98.

Ehrhardt-Martinez, K., Donnelly, K. A. and Laitner, S. (2010), Advanced metering initiatives and residential feedback programs: A meta-review for household electricity-saving opportunities, Technical report, American Council for an Energy Efficient Economy, Washington, D.C.

Fan, J. and Li, R. (2001), 'Variable selection via nonconcave penalized likelihood and its oracle properties', *Journal of the American Statistical Association* **96**(456), 1348–1360.

Farrar, D. E. and Glauber, R. R. (1967), 'Multicollinearity in regression analysis: The problem revisited', *The Review of Economics and Statistics* **49**(1), 92.

Filippín, C., Ricard, F. and Flores Larsen, S. (2013), 'Evaluation of heating energy consumption patterns in the residential building sector using stepwise selection and multivariate analysis', *Energy and Buildings* **66**, 571–581.

Fischer, C. (2008), 'Feedback on household electricity consumption: A tool for saving energy?', *Energy Efficiency* **1**(1), 79–104.

Fisher, R. J. (1993), 'Social desirability bias and the validity of indirect questioning', *Journal of Consumer Research* **20**(2), 303–315.

Friedman, J., Hastie, T. and Tibshirani, R. (2001), *The Elements of Statistical Learning*, Vol. 1 of *Springer Series in Statistics*, Springer, Berlin, Germany.

Friedman, J., Hastie, T. and Tibshirani, R. (2010), 'Regularization paths for generalized linear models via coordinate descent', *Journal of Statistical Software* **33**(1), 1.

Fumo, N. and Rafe Biswas, M. (2015), 'Regression analysis for prediction of residential energy consumption', *Renewable and Sustainable Energy Reviews* **47**, 332–343.

Gadenne, D., Sharma, B., Kerr, D. and Smith, T. (2011), 'The influence of consumers' environmental beliefs and attitudes on energy saving behaviours', *Energy Policy* **39**(12), 7684–7694.

Gardner, G. T. and Stern, P. C. (2009), 'The short list: The most effective actions U.S. households can take to curb climate change', *Environment Magazine* .

Gram-Hanssen, K. (2013), 'Efficient technologies or user behaviour, which is the more important when reducing households' energy consumption?', *Energy Efficiency* **6**(3), 447–457.

Gram-Hanssen, K. (2014), 'New needs for better understanding of household's energy consumption–behaviour, lifestyle or practices?', *Architectural Engineering and Design Management* **10**(1-2), 91–107.

Grubler, A., Wilson, C., Bento, N., Boza-Kiss, B., Krey, V., McCollum, D. L., Rao, N. D., Riahi, K., Rogelj, J., De Stercke, S., Cullen, J., Frank, S., Fricko, O., Guo, F., Gidden, M., Havlík, P., Huppmann, D., Kiesewetter, G., Rafaj, P., Schoepp, W. and Valin, H. (2018), 'A low energy demand scenario for meeting the 1.5°C target and sustainable development goals without negative emission technologies', *Nature Energy* **3**(6), 515–527.

Guerra-Santin, O. and Itard, L. (2010), 'Occupants' behaviour: Determinants and effects on residential heating consumption', *Building Research & Information* **38**(3), 318–338.

Guo, Y., Wang, J., Chen, H., Li, G., Liu, J., Xu, C., Huang, R. and Huang, Y. (2018), 'Machine learning-based thermal response time ahead energy demand prediction for building heating systems', *Applied Energy* **221**, 16–27.

Halvorsen, R. and Palmquist, R. (1980), 'The interpretation of dummy variables in semilogarithmic equations', *The American Economic Review* **70**(3), 474–475.

Harrell, F. (2001), *Regression Modeling Strategies: With Applications to Linear Models, Logistic Regression, and Survival Analysis*, Springer Series in Statistics, Springer-Verlag, New York, NY.

Heinze, G., Wallisch, C. and Dunkler, D. (2018), 'Variable selection–A review and recommendations for the practicing statistician', *Biometrical Journal* **60**(3), 431–449.

Hirst, E. and Goeltz, R. (1985), 'Comparison of actual energy savings with audit predictions for homes in the North Central Region of the U.S.A.', *Building and Environment* **20**(1), 1–6.

Hoerl, A. E. and Kennard, R. W. (1970), 'Ridge regression: Biased estimation for nonorthogonal problems', *Technometrics* **12**(1), 55.

Hoffman, I. M., Goldman, C. A., Murphy, S., Frick, N. A., Leventis, G. and Schwartz, L. C. (2018), The cost of saving electricity through energy efficiency programs funded by utility customers: 2009–2015, Technical Report 1457014, Lawrence Berkeley National Laboratory, Berkeley, CA.

Honaker, J. and King, G. (2010), 'What to do about missing values in time-series cross-section data', *American Journal of Political Science* **54**(2), 561–581.

Honaker, J., King, G. and Blackwell, M. (2011), 'Amelia II: A Program for Missing Data', *Journal of Statistical Software* **45**(1), 1–47.

Hsu, D. (2015), 'Identifying key variables and interactions in statistical models of building energy consumption using regularization', *Energy* **83**, 144–155.

Huang, S., Safiullah, H., Xiao, J., Hodge, B.-M. S., Hoffman, R., Soller, J., Jones, D., Dininger, D., Tyner, W. E., Liu, A. and Pekny, J. F. (2012), 'The effects of electric vehicles on residential households in the city of Indianapolis', *Energy Policy* **49**, 442–455.

Huebner, G. M., Cooper, J. and Jones, K. (2013), 'Domestic energy consumption—What role do comfort, habit, and knowledge about the heating system play?', *Energy and Buildings* **66**, 626–636.

Huebner, G., Shipworth, D., Hamilton, I., Chalabi, Z. and Oreszczyn, T. (2016), 'Understanding electricity consumption: A comparative contribution of building factors, socio-demographics, appliances, behaviours and attitudes', *Applied Energy* **177**, 692–702.

Jain R. K., Damoulas T. and Kontokosta C. E. (2014), Towards data-driven energy consumption forecasting of multi-family residential buildings: Feature selection via the lasso, *in* 'Proceedings of the 2014 International Conference on Computing in Civil and Building Engineering', American Society of Civil Engineers, Orlando, Florida, pp. 1675–1682.

Jain, R. K., Smith, K. M., Culligan, P. J. and Taylor, J. E. (2014), 'Forecasting energy consumption of multi-family residential buildings using support vector regression: Investigating the impact of temporal and spatial monitoring granularity on performance accuracy', *Applied Energy* **123**, 168–178.

Jones, R. V., Fuertes, A. and Lomas, K. J. (2015), 'The socio-economic, dwelling and appliance related factors affecting electricity consumption in domestic buildings', *Renewable and Sustainable Energy Reviews* **43**, 901–917.

Kavgic, M., Mavrogianni, A., Mumovic, D., Summerfield, A., Stevanovic, Z. and Djurovic-Petrovic, M. (2010), 'A review of bottom-up building stock models for energy consumption in the residential sector', *Building and Environment* **45**(7), 1683–1697.

Kavousian, A., Rajagopal, R. and Fischer, M. (2013), 'Determinants of residential electricity consumption: Using smart meter data to examine the effect of climate, building characteristics, appliance stock, and occupants' behavior', *Energy* **55**, 184–194.

Kelly, S. (2011), 'Do homes that are more energy efficient consume less energy?: A structural equation model of the English residential sector', *Energy* **36**(9), 5610–5620.

Kennedy, E. H., Beckley, T. M., McFarlane, B. L. and Nadeau, S. (2009), 'Why we don't "walk the talk": Understanding the environmental values/behaviour gap in Canada', *Human Ecology Review* **16**(2), 151.

Kialashaki, A. and Reisel, J. R. (2013), 'Modeling of the energy demand of the residential sector in the United States using regression models and artificial neural networks', *Applied Energy* **108**, 271–280.

King, G., Honaker, J., Joseph, A. and Scheve, K. (2001), 'Analyzing incomplete political science data: An alternative algorithm for multiple imputation', *American Political Science Review* **95**(1), 49–69.

Kuhn, M. (2008), 'Building Predictive Models in R Using the caret Package', *Journal of Statistical Software* **28**(1), 1–26.

Lam, J. C., Wan, K. K. W., Cheung, K. L. and Yang, L. (2008), 'Principal component analysis of electricity use in office buildings', *Energy and Buildings* **40**(5), 828–836.

LBNL (2004), Developing and testing low power mode measurement methods, Technical Report P500-04-057, Prepared for California Energy Commission Public Interest Energy Research Programme by Lawrence Berkeley National Laboratory, Berkeley, CA.

Lee, J. D., Sun, D. L., Sun, Y. and Taylor, J. E. (2016), 'Exact post-selection inference, with application to the lasso', *The Annals of Statistics* **44**(3), 907–927.

Lindén, A.-L., Carlsson-Kanyama, A. and Eriksson, B. (2006), 'Efficient and inefficient aspects of residential energy behaviour: What are the policy instruments for change?', *Energy Policy* **34**(14), 1918–1927.

Lockhart, R., Taylor, J., Tibshirani, R. J. and Tibshirani, R. (2014), 'A significance test for the lasso', *Annals of statistics* **42**(2), 413.

Mansouri, I., Newborough, M. and Probert, D. (1996), 'Energy consumption in U.K. households: Impact of domestic electrical appliances', *Applied Energy* **54**(3), 211–285.

McLoughlin, F., Duffy, A. and Conlon, M. (2012), 'Characterising domestic electricity consumption patterns by dwelling and occupant socio-economic variables: An Irish case study', *Energy and Buildings* **48**, 240–248.

Meinshausen, N. and Bühlmann, P. (2006), 'High-dimensional graphs and variable selection with the Lasso', *The Annals of Statistics* **34**(3), 1436–1462.

Mundaca, L., Ürge-Vorsatz, D. and Wilson, C. (2018), 'Demand-side approaches for limiting global warming to 1.5°C', *Energy Efficiency* .

Ndiaye, D. and Gabriel, K. (2011), 'Principal component analysis of the electricity consumption in residential dwellings', *Energy and Buildings* **43**(2-3), 446–453.

NEETF & RoperASW (2002), Americans' low "Energy IQ": A risk to our energy future. Why America needs a refresher course on energy, Technical report, The National Environmental Education & Training Foundation & Roper ASW, Washington, D.C.

Nejat, P., Jomehzadeh, F., Taheri, M. M., Gohari, M. and Abd. Majid, M. Z. (2015), 'A global review of energy consumption, CO2 emissions and policy in the residential sector (with an overview of the top ten CO2 emitting countries)', *Renewable and Sustainable Energy Reviews* **43**, 843–862.

Pampaka, M., Hutcheson, G. and Williams, J. (2016), 'Handling missing data: Analysis of a challenging data set using multiple imputation', *International Journal of Research & Method in Education* **39**(1), 19–37.

Pao, H. T. (2009), 'Forecasting energy consumption in Taiwan using hybrid nonlinear models', *Energy* **34**(10), 1438–1446.

Pisillo, J. (2012), 'Palo Alto residents earn 3rd highest median family income', `http://blog.sfgate.com/ontheblock/2012/10/12/palo-alto-residents-earn-3rd-highest-median-family-income/`.

R Core Team (2017), 'R: A language and environment for statistical computing', R Foundation for Statistical Computing.

Ratner, B. (2010), 'Variable selection methods in regression: Ignorable problem, outing notable solution', *Journal of Targeting, Measurement and Analysis for Marketing* **18**(1), 65–75.

Ripley, B. D. and Venables, W. N. (2002), *Modern Applied Statistics with S*, 4 edn, Springer, New York, NY.

Roberts, N. and Thatcher, J. (2009), 'Conceptualizing and testing formative constructs: Tutorial and annotated example', *ACM SIGMIS Database* **40**(3), 9–39.

Robinson, C., Dilkina, B., Hubbs, J., Zhang, W., Guhathakurta, S., Brown, M. A. and Pendyala, R. M. (2017), 'Machine learning approaches for estimating commercial building energy consumption', *Applied Energy* **208**, 889–904.

Rubin, D. B. (1996), 'Multiple imputation after 18+ years', *Journal of the American Statistical Association* **91**(434), 473.

Sanquist, T. F., Orr, H., Shui, B. and Bittner, A. C. (2012), 'Lifestyle factors in U.S. residential electricity consumption', *Energy Policy* **42**, 354–364.

Schultz, P. W., Estrada, M., Schmitt, J., Sokoloski, R. and Silva-Send, N. (2015), 'Using in-home displays to provide smart meter feedback about household electricity consumption: A randomized control trial comparing kilowatts, cost, and social norms', *Energy* **90**, 351–358.

Shmueli, G. (2010), 'To explain or to predict?', *Statistical Science* **25**(3), 289–310.

Sidhu, I., Kaminsky, P., Tenderich, B., DeForest, N., Lorimer, A. and Ur, B. (2009), Impact of widespread electric vehicle adoption on the electrical utility business–threats and opportunities, Technical report, University of California, Berkeley Center for Entrepreneurship & Technology (CET), Berkeley, CA.

Sorrell, S., Dimitropoulos, J. and Sommerville, M. (2009), 'Empirical estimates of the direct rebound effect: A review', *Energy Policy* **37**(4), 1356–1371.

Southwell, B., Murphy, J., DeWaters, J. and LeBaron, P. (2012), Americans' perceived and actual understanding of energy, Technical report, RTI Press, Research Triangle Park, NC.

Steemers, K. and Yun, G. Y. (2009), 'Household energy consumption: A study of the role of occupants', *Building Research & Information* **37**(5-6), 625–637.

Suryanarayana, G., Lago, J., Geysen, D., Aleksiejuk, P. and Johansson, C. (2018), 'Thermal load forecasting in district heating networks using deep learning and advanced feature selection methods', *Energy* **157**, 141–149.

Swan, L. G. and Ugursal, V. I. (2009), 'Modeling of end-use energy consumption in the residential sector: A review of modeling techniques', *Renewable and Sustainable Energy Reviews* **13**(8), 1819–1835.

Takahashi, M. (2017), 'Statistical inference in missing data by MCMC and non-MCMC multiple imputation algorithms: Assessing the effects of between-imputation iterations', *Data Science Journal* **16**(0), 37.

Taylor, J. and Tibshirani, R. J. (2015), 'Statistical learning and selective inference', *Proceedings of the National Academy of Sciences* **112**(25), 7629–7634.

Thøgersen, J. and Grønhøj, A. (2010), 'Electricity saving in households—A social cognitive approach', *Energy Policy* **38**(12), 7732–7743.

Tibshirani, R. (1996), 'Regression shrinkage and selection via the lasso', *Journal of the Royal Statistical Society. Series B (Methodological)* pp. 267–288.

Tiwari, P. (2000), 'Architectural, demographic, and economic causes of electricity consumption in Bombay', *Journal of Policy Modeling* **22**(1), 81–98.

Tso, G. K. and Guan, J. (2014), 'A multilevel regression approach to understand effects of environment indicators and household features on residential energy consumption', *Energy* **66**, 722–731.

Tso, G. K. and Yau, K. K. (2007), 'Predicting electricity energy consumption: A comparison of regression analysis, decision tree and neural networks', *Energy* **32**(9), 1761–1768.

Uniejewski, B., Nowotarski, J. and Weron, R. (2016), 'Automated variable selection and shrinkage for day-ahead electricity price forecasting', *Energies* **9**(8), 621.

U.S. Census Bureau (2015), '2011-2015 American Community Survey', `https://www.census.gov/acs/www/data/data-tables-and-tools/`.

U.S. DOE (2014), Evaluating electric vehicle charging impacts and customer charging behaviors–experiences from six smart grid investment grant projects, Technical report, U.S. Department of Energy Office of Electricity Delivery & Energy Reliability, Washington, D.C.

U.S. EIA (2016), 'How much electricity does an American home use?', `https://www.eia.gov/tools/faqs/faq.php?id=97&t=3`.

U.S. EIA (2018), Annual Energy Outlook 2018 with projections to 2050, Technical report, U.S. Energy Information Administration, Washington, D.C.

U.S. EPA (2015), Regulatory impact analysis for the Clean Power Plan final rule, Technical Report EPA-452/R-15-003, U.S. Environmental Protection Agency, Research Triangle Park, NC.

U.S. EPA (2017), Inventory of U.S. greenhouse gas emissions and sinks: 1990-2015, Technical report, U.S. Environmental Protection Agency, Washington, D.C.

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* **16**(1), 98–105.

Vringer, K., Aalbers, T. and Blok, K. (2007), 'Household energy requirement and value patterns', *Energy Policy* **35**(1), 553–566.

Wahlström, M. H. and Hårsman, B. (2015), 'Residential energy consumption and conservation', *Energy and Buildings* **102**, 58–66.

Wallis, H., Nachreiner, M. and Matthies, E. (2016), 'Adolescents and electricity consumption: Investigating sociodemographic, economic, and behavioural influences on electricity consumption in households', *Energy Policy* **94**, 224–234.

Wang, M., Wright, J., Brownlee, A. and Buswell, R. (2016), 'A comparison of approaches to stepwise regression on variables sensitivities in building simulation and analysis', *Energy and Buildings* **127**, 313–326.

Wang, Z., Wang, Y., Zeng, R., Srinivasan, R. S. and Ahrentzen, S. (2018), 'Random forest based hourly building energy prediction', *Energy and Buildings* **171**, 11–25.

Wei, Y., Zhang, X., Shi, Y., Xia, L., Pan, S., Wu, J., Han, M. and Zhao, X. (2018), 'A review of data-driven approaches for prediction and classification of building energy consumption', *Renewable and Sustainable Energy Reviews* **82**, 1027–1047.

Wiesmann, D., Lima Azevedo, I., Ferrão, P. and Fernández, J. E. (2011), 'Residential electricity consumption in Portugal: Findings from top-down and bottom-up models', *Energy Policy* **39**(5), 2772–2779.

Wilhite, H., Shove, E., Lutzenhiser, L. and Kempton, W. (2000), Twenty years of energy demand management: We know more about individual behavior but how much do we really know about demand, *in* 'Proceedings of the 2000 ACEEE Summer Study on Energy Efficiency in Buildings', American Council for an Energy Efficient Economy, Pacific Grove, CA, pp. 435–453.

WRCC (2006), 'Palo Alto, California: Period of record general climate summary 1953-2006', `https://wrcc.dri.edu/cgi-bin/cliMAIN.pl?capalo+sfo`.

Wyatt, P. (2013), 'A dwelling-level investigation into the physical and socio-economic drivers of domestic energy consumption in England', *Energy Policy* **60**, 540–549.

Yang, Y. and Zou, H. (2017), 'Group lasso penalized learning using a unified BMD algorithm'.

Yohanis, Y. G., Mondol, J. D., Wright, A. and Norton, B. (2008), 'Real-life energy use in the U.K.: How occupancy and dwelling characteristics affect domestic electricity use', *Energy and Buildings* **40**(6), 1053–1059.

Yu, Z., Haghighat, F., Fung, B. C. M. and Yoshino, H. (2010), 'A decision tree method for building energy demand modeling', *Energy and Buildings* **42**(10), 1637–1646.

Yuan, M. and Lin, Y. (2006), 'Model selection and estimation in regression with grouped variables', *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **68**(1), 49–67.

Zhao, H.-X. and Magoulès, F. (2010), 'Parallel support vector machines applied to the prediction of multiple buildings energy consumption', *Journal of Algorithms & Computational Technology* **4**(2), 231–249.

Zhao, H.-X. and Magoulès, F. (2012), 'A review on the prediction of building energy consumption', *Renewable and Sustainable Energy Reviews* **16**(6), 3586–3592.

Zhao, P. and Yu, B. (2006), 'On model selection consistency of lasso', *Journal of Machine Learning Research* **7**, 2541–2563.

Zou, H. (2006), 'The adaptive lasso and its oracle properties', *Journal of the American Statistical Association* **101**(476), 1418–1429.

Zou, H. and Hastie, T. (2005), 'Regularization and variable selection via the elastic net', *Journal of the Royal Statistical Society. Series B (Statistical Methodology)* **67**(2), 301–320.

4.A Supplemental figures

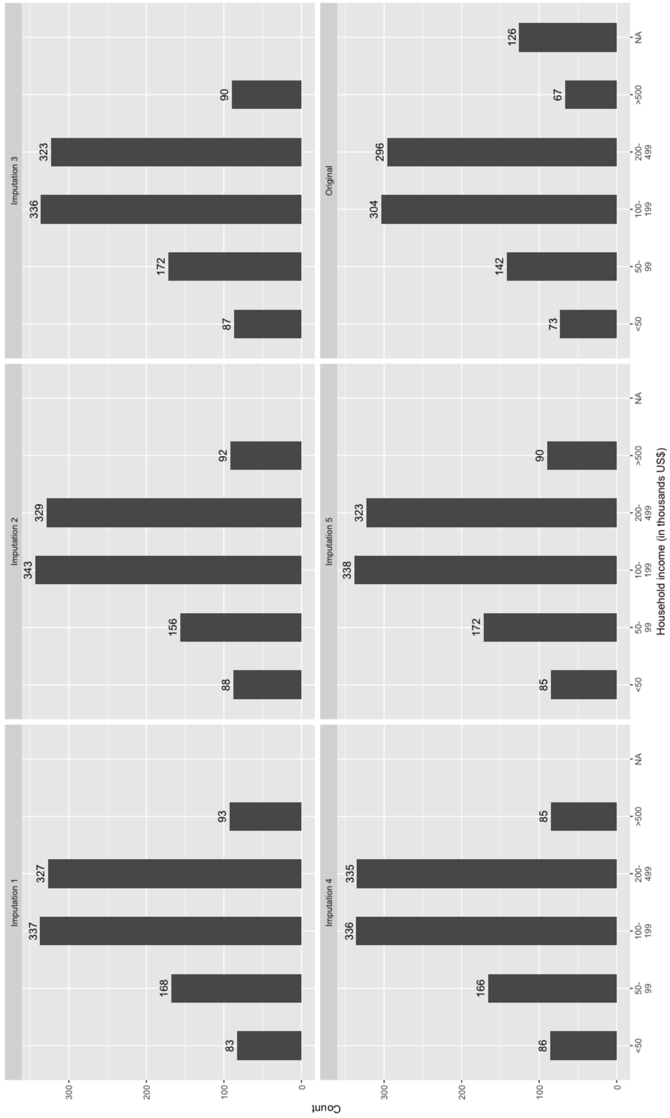

Figure 4.A.1: Distributions of *Income* variable across imputed and original data sets.

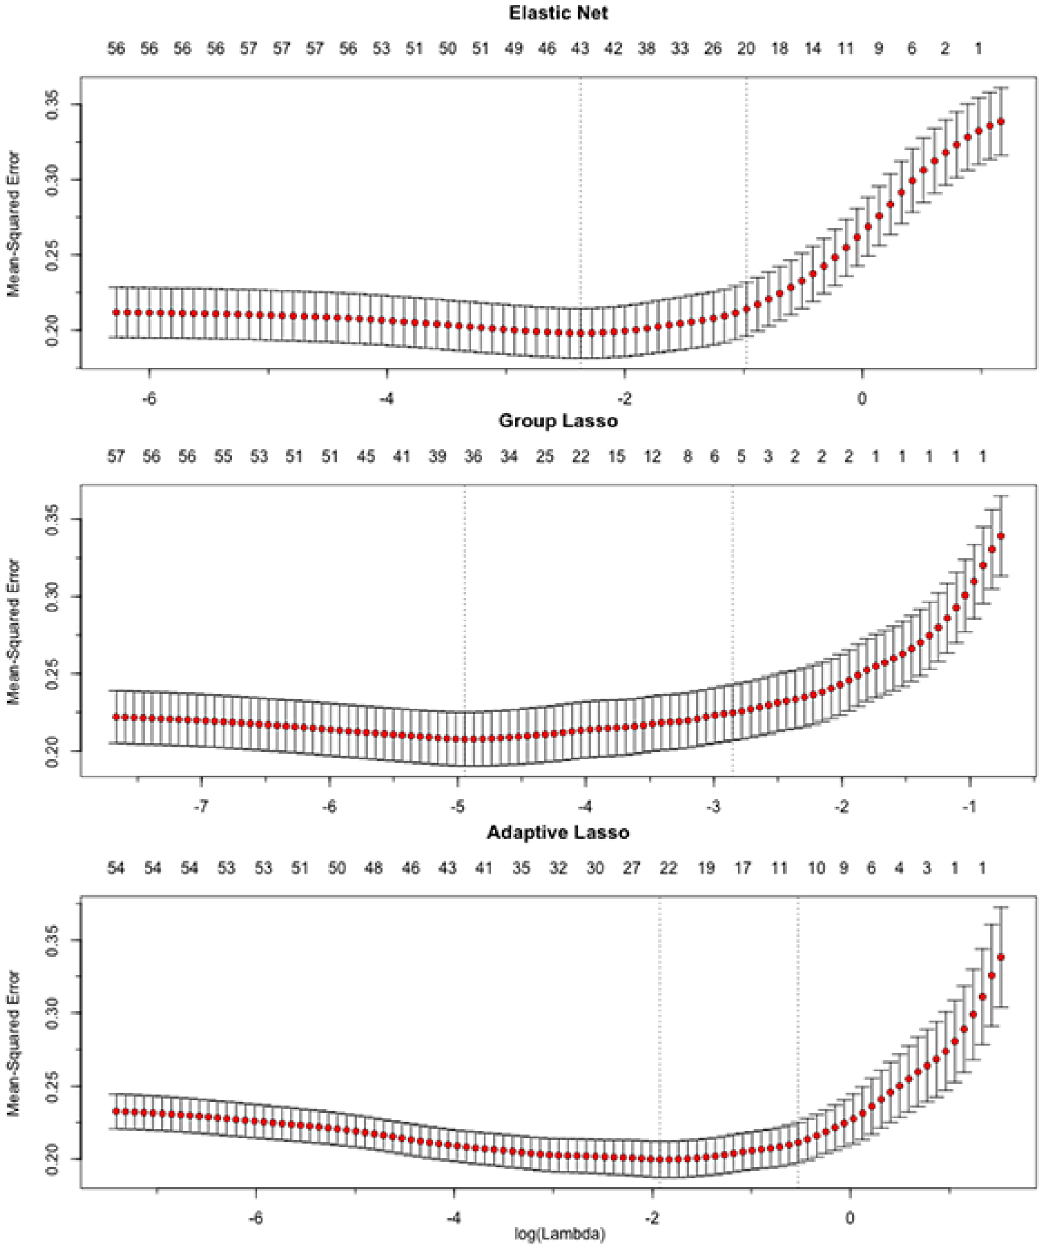

Figure 4.A.2: Plots of cross-validation MSE for elastic net, group lasso, and adaptive lasso applied to the data filtered for AC ownership ($N = 390$). Axes and plots elements are the same as in Figures 4.2–4.3.

Chapter 5

Daily life and demand: An analysis of intra-day variations in residential electricity demand with time-use data (Paper 2)

Aven Satre-Meloy, Marina Diakonova, Philipp Grünewald
Published in *Energy Efficiency*

Abstract

Demand-side flexibility has been suggested as a tool for peak demand reduction and large-scale integration of low-carbon electricity sources. Deeper insight into the activities and energy services performed in households could help to understand the scope and limitations of demand-side flexibility. Measuring and Evaluating Time- and Energy-use Relationships (METER) is a five-year, U.K.-based research project and the first study to collect activity data and electricity use in parallel at this scale. We present statistical analyses of these new data, including more than 6,250 activities reported by 450 individuals in 173 households, and their relationship to electricity use patterns. We use a regularisation technique to select influential variables in regression models of average electricity use over a day and of discretionary use across four-hour time periods to compare intra-day variations. We find that dwelling and appliance variables are the strongest predictors and can explain 49% of the variance in mean daily demand. For models of four-hour average 'de-minned' demand, we find that activity variables are consistently influential, both in terms of coefficient magnitudes and contributions to increased model explanatory power. Activities relating to food preparation and eating, household chores, and recreation show the strongest associations. We conclude that occupant activity data can advance our understanding of the temporal characteristics

of electricity demand and inform approaches to shift or reduce it. We stress the importance of considering demand as a function of time of day, and we use our findings to argue that a more nuanced understanding of this relationship can yield useful insights for residential demand flexibility.

5.1 Introduction

The residential sector is the largest end user of electricity in the U.K., accounting for 30% of total consumption in 2017 (BEIS, 2018). It is also responsible for up to 50% of national peak demand, during which time electricity provision is especially costly and carbon-intensive (Ofgem, 2010). As evidenced by the share of generation from renewables increasing to over 29% in 2017, the U.K.'s electricity system is increasingly low-carbon, yet much more ambition is needed to achieve national climate goals in the next decades (BEIS, 2018).

The U.K.'s 2017 Clean Growth Strategy sets forth a set of policies and proposals to further accelerate the deployment of low-carbon energy while maintaining increased economic growth (BEIS, 2017). One of the proposals, the 'Smart systems plan', aims to help consumers use energy more flexibly. Energy storage will support integration efforts as the U.K. pursues rapid deployment of renewable energy, but demand flexibility can reduce the cost of integration and storage requirements for low-carbon energy systems (National Infrastructure Commission, 2016).

Understanding how to make demand more flexible in the residential building sector requires a more detailed investigation of what electricity is used for in households (Grunewald and Diakonova, 2018*b*; Powells et al., 2014). Time-use data are increasingly used for modelling electricity consumption, but these analyses often use simulated activity data and proxies for electricity consumption, such as household expenditure data, which require assumptions about the link between time-use data and electricity consumption patterns. In some cases, these models are validated against real consumption data, though the validation sample sizes are in the 10–20 households range (Richardson et al., 2010; Widén et al., 2009;

Widén and Wäckelgård, 2010). Even as time-use data have been more recently incorporated into these models, however, there remains a notable gap in the literature of empirical studies on how occupant activities influence the temporal aspects of demand (Anderson, 2016; Anderson and Torriti, 2018; Grunewald and Diakonova, 2018*a*).

In this paper, we aim to contribute to filling this gap by constructing regression models of original survey, time-use, and electricity demand data for 173 U.K. households collected as part of an ongoing study. We test the extent to which socio-demographic, dwelling, appliance, and activity data can explain variations in intra-day household electricity use, from 5 p.m. until 9 p.m. on the following day. We then compare the effects of reported time-use activities in households throughout the day by dividing the day into six equal four-hour periods.

We aim to understand whether and how the inclusion of categorised activity data improve models of household electricity demand. We investigate how different categories of activities are associated with electricity demand during different times of day, and we consider how this knowledge can inform efforts to enhance demand-side flexibility and enable more and less costly integration of renewable energy sources.

The paper is organised as follows: In Section 5.2, we present a brief review of the flexibility and household electricity modelling literature, including both determinants of household electricity consumption and the use of time-use data in energy modelling. In Section 5.3 we present our methodology for collecting household electricity and time-use data in parallel. Section 5.4 presents results of full-day and four-hour models of average electricity demand. In Section 5.5, we discuss the implications of our findings, and in Section 5.6 we conclude.

5.2 Literature review

5.2.1 Demand-side flexibility

Flexibility has been defined in the energy demand literature as the ability for consumers to change how, when, and where energy is used—a shift in focus from the magnitude of overall consumption to the timing of demand (Powells et al., 2014; Torriti et al., 2015). We conceive of it here simply as "a potential to change power at a certain rate (Watt/hour)" (Grunewald and Diakonova, 2018*b*, p. 59). While much of the focus in demand-side energy research has been on reducing demand through energy efficiency, as the energy system becomes increasingly supplied by variable, low-carbon generation, flexibility becomes increasingly important for balancing supply and demand. The potential responsiveness of demand, or its ability to shift in time to match high generation from renewables or to flatten demand during peak periods, is essential for minimising the costs of transitioning to a low-carbon power system (Strbac et al., 2012).

Demand-side response (DSR) refers to measures that provide flexibility to the energy system by shifting or reshaping energy loads. McKenna et al. (2017) conceive of DSR as a three-dimensional space including 'technology change', 'service expectation change', and 'activity change'. The technology change dimension refers to an automated or remotely controlled shift in energy use, for example a smart fridge that delivers a response without affecting the energy service delivered. The authors note this has been the main form of DSR that has been tested in energy models given the relative ease of simulating such technology changes. 'Service expectation change' refers to shifting occupants' expected level of energy service from an appliance or technology (e.g., the thermostat set-point), while an 'activity change' requires a shift in either the timing or type of activity and relies on more active behavioural changes among energy users.

Grunewald and Diakonova (2018*b*) more broadly differentiate these dimensions of DSR as 'appliance led' and 'activity led'. Among 'activity led' DSR are several different types of shift mechanisms: activity shifts, or changing the timing of an activity; substituting practice, which does not reorder activities in time but instead substitutes an energy-intensive activity for one that is not (e.g., having a cold rather than hot meal); substituting metabolic energy, such as mixing dough by hand rather than with an electric mixer; or changing the practitioner, which might entail going out for dinner rather than cooking at home.

Recent research has begun to investigate the DSR potential of these and similar types of measures, but still little is known about both the mechanisms by which electricity consumption patterns can be reshaped and the different constraints or motivating factors that can reshape them. A review of major DSR trials in the U.K. found that residential customers are responsive to economic incentives to shift demand but that the size of the shift can vary significantly (DECC, 2012). The report notes several areas where findings were inconclusive. These include the responsiveness of vulnerable and low-income customers, the impact of non-economic signals, and detailed information on the way consumers shift their demand in response to incentives.

Further empirical work on the flexibility of household activities includes research that qualitatively assessed the likelihood of certain practices being performed during peak hours (4 to 8 p.m.) (Powells et al., 2014). This study found a very high or high likelihood of TV watching, cooking, computer/Internet use, and dish washing during these hours, with lesser frequencies of laundry, ironing, vacuuming, and bathing. More insight into the types of activities happening during peak periods and their contributions to electricity demand can help identify those that are best suited for DSR.

But how well 'suited' an activity is for shifting is not the same as how 'flexible' it might be. Torriti et al. (2015) stress this point by developing a 'Flexibility Index' composed of five separate indices (synchronisation, variation, non-shared activities, active home occupancy, and spatial mobility). These component indices are calculated for five time periods throughout the day for a sample of 153 respondents, and the findings show that morning peak times feature high levels of synchronisation and shared activities but low variation in activities. Evening peak times show less synchronisation but higher spatial mobility and variation in activities.

Demand-side flexibility is complex and is difficult to measure empirically. It is likely that socio-demographic, dwelling, appliance ownership characteristics, and activity patterns bear a strong relationship to both the types of DSR that can take place in homes as well as the capacity of households to participate in DSR programs. Several different approaches have been developed to better understand these relationships in order to inform more effective DSR. The following sections review these approaches in the context of research on electricity demand modelling and incorporating time-use data in these models.

5.2.2 Determinants of household electricity consumption

The literature on quantitative approaches to model and understand key socio-technical determinants of electricity use in households is expansive. Previous studies employ both top-down and bottom-up approaches. Top-down approaches consider how national-level data, such as characteristics of the housing stock and macroeconomic factors, influence electricity consumption, while bottom-up approaches use detailed household-level data on dwelling characteristics and occupant socio-demographics to identify the primary drivers of electricity consumption for a sample of households and then extrapolate these findings to the wider population (Swan and Ugursal, 2009). Studies are also varied in how they measure household electricity use, in some cases considering average annual consumption values while in others studying daily or hourly consumption.

A recent comprehensive review of the socio-demographic, dwelling, and appliance related determinants of household electricity consumption undertaken by Jones et al. (2015) shows that no fewer than 62 factors have been found to affect household electricity consumption. Twenty of these are consistently found to have a significant, positive correlation with electricity use ('consistently' is defined by the authors as showing a statistically significant relationship in more than three studies). The socio-demographic factors include number of occupants, presence of teenagers, household income, and disposable income. The impact of occupant age, education level, and tenure type is inconclusive in the studies reviewed. The dwelling factors that are influential include dwelling age, dwelling type, number of rooms, total floor area, and ownership of electric space heating and cooling systems. Appliance factors include ownership of a desktop computer, television, electric oven, refrigerator, dishwasher, washing machine, and tumble dryer, as well as the overall number of appliances owned.

Huebner et al. (2016) confirm many of these findings in one of the few representative studies of electricity consumption in the U.K. residential sector. Using data from a sample of 845 English households from 2011 to 2012, they find that models containing only appliance ownership and use factors explain 34% of the variance in non-heating annual electricity consumption, whereas models containing socio-demographic variables explain only 21%, and models containing dwelling and other occupant variables are poor in explaining electricity use. Their model combining these factors shows that dwelling floor area, number of occupants, wet appliance ownership, and hours of TV watched per day are statistically significant predictors of annual consumption.

As our paper focuses specifically on factors affecting intra-day variations in electricity use, a brief review of findings from studies with a similar focus is included here. The number of these studies is fewer than those included in the review discussed above because in much of the previous research, advanced metering data were not as available or accessible. With the growth of smart meter adoption, however, these data and studies investigating them are becoming increasingly common.

In general, this body of work suggests that appliance ownership and usage factors have stronger relationships with variations in consumption than socio-demographic and dwelling factors. In a study monitoring 27 U.K. residential buildings' electricity consumption at five-minute intervals for a period of two years, Firth et al. (2008) find that an observed increase in usage between years is primarily attributable to increases in the consumption of standby appliances, such as televisions and other consumer electronics, and 'active' appliances, such as kettles, washing machines, electric showers, and lighting.

McLoughlin et al. (2012) study the effect of dwelling and occupant characteristics for a representative sample of 4,200 Irish dwellings on several dependent variables, including total electricity consumption over a six-month period, mean daily maximum demand over that period, electrical load factor (ratio of daily mean to daily maximum electrical demand), and timing of peak demand. They find that several dwelling and socio-demographic factors, such as number of bedrooms, type of dwelling, and head of household age, are significant in explaining load factor variations (i.e., those that are 'peakier' or 'flatter'). They also find that timing of peak demand is mostly driven by variables measuring reported appliance ownership and use, especially dishwashers, electric power showers, televisions, and desktop computers. Notably, the models are less able to explain load factor and timing of peak demand than total consumption and mean daily maximum demand. Key explanatory factors for understanding variations in electricity consumption at different times of day may therefore be missing in the models.

Exploiting a similar smart meter data set from the Irish Commission for Energy Regulation (CER), Anderson et al. (2016) examine the links between half-hourly electricity consumption data and household characteristics for the purpose of assessing the feasibility of using high-resolution electricity data to infer household characteristics to supplement census-taking efforts. They construct multi-level regression models to analyse numerous 'profile indicators' that describe the temporal characteristics of load profiles, such as base load, mean load, 97.5th percentile load (as a proxy for peak load), and several other temporal parameters. Their results show that several household characteristics, such as number of occupants, household income, and presence of children, are useful predictors of several of these 'profile indicators', especially those measuring overall magnitude of consumption.

Kavousian et al. (2013) examine structural and behavioural determinants of daily maximum and minimum demand for a data set of 10-minute interval electricity readings and an extensive household survey of 1,628 U.S. dwellings in California. The study was non-random and had a sample biased toward high-income and well-educated participants. They find that while daily minimum demand is influenced most by weather, location, and physical characteristics of the building, such as size and type of home, daily maximum demand is influenced by high-consumption, intermittent appliances, such as electric water heaters or tumble dryers. The authors note the clear difference in their results between drivers of daily minimum and maximum demand, explaining that peak demand is more dependent on the activity patterns that lead to use of high-consumption appliances. Minimum demand is driven primarily by locational and physical characteristics of the dwelling.

Several recent papers have aimed to investigate socio-demographic and physical dwelling determinants of whole load profile shapes rather than average values of electricity consumption across varying time periods (McLoughlin et al., 2015; Rhodes et al., 2014; Viegas et al., 2016). These studies take an alternative approach by first clustering load profiles into representative profile classes and then using logistic regression to determine factors influencing membership in different load profile classes. While this is a different approach than the one used in this paper and in the studies reviewed above, it achieves a similar aim. The findings from these studies indicate that variables such as working from home, time spent watching television, age, and ownership of dishwashers and washing machines are most important for explaining the shape of a household's daily electrical load profile.

The literature investigating highly-resolved variations in residential electricity consumption is growing alongside the availability of data, but many of the results discussed here are dependent on the sample nature and size and not necessarily generalisable to wider populations. It is also important to consider the climatic and social contexts in which these studies are undertaken, as these likely influence their findings.

5.2.3 Incorporating time-use data in studies of electricity consumption

Social scientists have long researched the dynamics of daily life through studies of how people spend their time. Longitudinal studies of thousands of households have investigated questions of historical changes in time use and how these relate to trends in technological and social change (eurostat, 2009; Gershuny, 2008; U.S. Bureau of Labor Statistics, 2018). Research in energy demand modelling has acknowledged the need to take account of the timing of household activities in order to more accurately represent electricity demand

(Torriti et al., 2015). In the last decade, occupant activity data from time-use studies have been increasingly incorporated into high-resolution models of residential electricity demand (Torriti, 2014). These studies involve varying approaches for incorporating time-use data. In some studies, probabilistic simulations are used to develop models of active occupancy, occupant activity sequences, and appliance usage that are predictive of electricity load curves (Ellegård and Palm, 2011; Richardson et al., 2010; Widén et al., 2009). In others, changes in electricity consumption over decades are linked to changes in time-use and household expenditures data, often using decomposition analysis, which attributes changes in these factors to overall changes in electricity use (De Lauretis et al., 2017; Jalas and Juntunen, 2015; Sekar et al., 2018). In many of these studies, models and results are validated with actual measurements from a small sample of households (McKenna and Thomson, 2016; Widén et al., 2009).

While this literature shows varying results for the links between time use and electricity consumption, some common themes emerge. Mealtimes and related activities are consistently found to be high consumption activities (De Lauretis et al., 2017; Druckman et al., 2012; Jalas and Juntunen, 2015). Studies also find evidence that housework and personal time are energy-intensive activities (Palmer et al., 2013; Widén et al., 2009). Unsurprisingly, sleeping and resting are low-consumption activities (Druckman et al., 2012; Jalas and Juntunen, 2015), whereas recreational activities can be either high or low energy depending on whether they involve televisions and computers or reading and socialising. In terms of the temporal shifts in the timing of U.K. peak demand over the past 40 years, Anderson and Torriti (2018) find that this is mostly due to changes in food related activities, in part due to changes in personal care and housework shifts, and little to do with changes in media use.

Using the specific case of laundry practices in the U.K., Anderson (2016) highlights the role of societal change in time shifts for energy-using activities. Increases in labour market participation by more females in society is likely partially responsible for shifting laundry energy demand into weekday mornings and evening peak times but also into Sunday mornings, which are less problematic for providing energy. The author concludes that more analyses of the changes to routine energy-using practices are necessary for assessing the value of various DSR strategies.

A review of time-use modelling of electricity demand undertaken by Torriti (2014) identifies five limitations in the literature. First, time-use data must be aggregated across large numbers of households to be statistically significant. Second, time-use data are often sampled throughout the year, whereas peak electricity events often happen on specific days, such as during temperature extremes or national events. Third, the low frequency of large time-use surveys means the data quickly become outdated. Fourth, the timing of occupancy and typical activities in developed countries is less variable than other factors that influence consumption, such as weather or appliance design. Fifth, modelling multiple-person households using time-use data, especially for modelling occupancy patterns, is much more challenging than for single-person households, which was confirmed empirically by Grunewald and Diakonova (2018*a*). The fourth point, however, is contested. Tabbone et al. (2016) found that even within countries, cultural differences in time use can affect load profiles.

To these limitations, McKenna et al. (2017) add several additional failings of time-use demand models. First, they note that time-use studies were not originally designed for the purpose of modelling energy use, and they thus do not differentiate between 'energy-intensive' or 'low-energy' alternatives of the same activity. Meal time is a good example, as this can vary from relatively low energy (cold meal preparation) to very energy intensive (oven and electric hob use). Second, time-use data do not account for overlapping or 'bundled' activities, such as those that might occur during multi-tasking. Third, the typical 10-minute reporting window for activities may not capture energy-relevant yet shorter duration activities, such as boiling the kettle.

This study addresses the limitations highlighted in these two reviews through a new approach to collecting time-use data. To the authors' knowledge, this is the first study to collect time-use data and electricity readings directly from households in parallel at this scale. This research contributes to an emerging focus within energy demand studies on not only the magnitude but also the timing of electricity consumption in households and how it is influenced both by numerous socio-demographic and dwelling factors that have been studied in previous research, as well as by the activities of occupants, which may contribute additional insights for determining the timing of electricity use in homes and the potential for shifting these activities to provide more flexibility to the energy system.

5.3 Data and methodology

5.3.1 Data collection

Electricity readings and activity records are collected from U.K. households as part of an ongoing study (Grünewald and Layberry, 2015). Participation is voluntary, and participants are recruited online, via radio, and through campaigns at selected community events. An incentive for participation is the chance to win a year's worth of free electricity.

When registering for the study, participants complete a survey wherein they provide detailed socio-demographic information along with data on dwelling characteristics and appliance ownership. Participating households are sent a parcel prior to their chosen date, which includes an electricity recorder, activity recorder(s), and an instruction booklet. The study encourages all household members over the age of eight to participate.

Participants are instructed to attach the electricity recorder below the household's electricity meter. The electricity recorder collects readings every second for 28 hours, from 5 p.m. on their chosen day until 9 p.m. the following day. This study length is chosen in order to capture two typical peak demand periods and because the electricity recorders have a battery life of around 28 hours per charge. Fuel type used for heating and cooking appliances is collected in the household survey.

Activities are recorded with a dedicated app pre-installed on individual devices. The app presents six options per screen that guide participants through recording their activities, starting with the location and concluding with the number of people partaking in the activity and one's enjoyment of it. Asking about enjoyment was found to increase participant retention in previous time-use research (Gershuny and Sullivan, 2017). Figure 5.1 shows an example activity entry sequence. In contrast to paper-based time-use diaries (eurostat, 2009), the app can guide participants to select energy-relevant details, such as appliance use, allowing for detailed descriptions of activities. Instead of capturing the duration of activities, as was done in previous time-use studies, the app records activities reported instantaneously. Users can record multiple activities in sequence and are also given an option to record the 'end' of activities. While users are encouraged to report activities when they actually occur, entries can be made retrospectively and into the future. One test of validity for reporting accuracy suggests around 80% of activity records for boiling the kettle are reported within 10 minutes of the activity itself (see Grunewald and Diakonova (2018*a*) for more details). More details about the functionality of the app are discussed in Grünewald et al. (2017), and a description of the data storage and handling procedures can be found in Grünewald and Diakonova (2019).

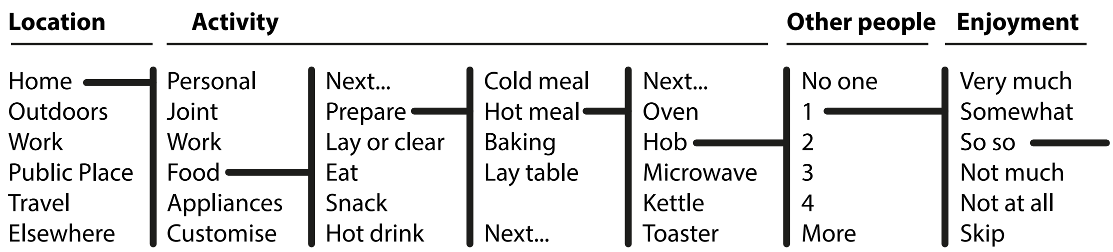

Figure 5.1: Example of an activity entrance sequence on the activity recorder (Grünewald et al., 2017).

5.3.2 Study sample

The study sample consists of 173 households and 447 individuals who together reported 6,265 activities. The process for determining inclusion in subsequent analyses is as follows:

- Households that own solar PV are excluded due to concerns about the validity of electricity readings.

- Households that did not complete the full survey are excluded. The small occurrence of missing data (23 cases) did not warrant more complex handling, such as multiple imputation, so these cases are deleted listwise.

- Electricity readings below 20 W are excluded from daily and four-hour averages, as this signals a failure to attach the electricity recorder properly.

- Activities that were reported either before the electricity meter starts recording or after it stops are excluded. Similarly, activities that are reported in an hour where a valid electricity reading is missing are excluded.

- Only activities reported in the home are included, since we are interested in those activities that have a direct influence on household electricity demand.

5.3.3 Socio-demographic, dwelling, and appliance variables

The household survey includes questions on occupant socio-demographics, physical dwelling characteristics, and appliance ownership. In previous studies, these have shown to have the strongest associations with residential electricity consumption in the U.K. (Huebner et al., 2016; Jones and Lomas, 2016; Wyatt, 2013). Table 5.1 presents frequencies and descriptive statistics for these variables for both the sample and the wider U.K. population (where available). Categorical variables are dummy coded prior to analysis, and reference categories are bolded in Table 5.1. Estimated monthly electric bill measures participants' estimate of their electricity expenditure. While we would expect this to have a positive association to average daily electricity demand, we suspect that monthly estimates on expenditure might vary widely from day-to-day variations in demand. We include it in our analyses to understand how well occupant estimates of electricity demand do in fact predict actual average daily demand values.

The study is voluntary, and the sample is not fully representative of the general U.K. population. Selection biases include an underrepresentation of renters, low-income groups, and the elderly. Gas boiler ownership is slightly underrepresented, suggesting a higher prevalence of all-electric heating in the sample. These biases are likely to influence both overall consumption as well as the timing of demand, which is why we caution against overgeneralising from these results.

Table 5.1 also lists the distributions of season and day of week that households participate. A majority of households (57%) participate during the winter season (November–February). Data collection is focused on this season because it coincides with U.K. annual peak demand (DECC, 2014). Households are recruited to start the study on a weekday, though given the study period spans two days, 22% of the sample starts on a Friday and completes the study on a Saturday. We do not create separate models for season or day of week in order to preserve the sample size across models, but we do include these as candidate variables during model selection, which is further explained below.

5.3.4 Activity variables

The app pre-codes activities, following the Harmonised European Time-Use Studies (HETUS) standards. For this analysis, activities are grouped into seven categories, which have become common in time-use research (Ellegård et al., 2010; Lader et al., 2006; Stankovic et al., 2016). Examples of activities that fall into each category are shown in Table 5.2. In the following analyses, activities are aggregated across household members so that each household has a total count, or frequency, for each activity category. Table 5.2 presents descriptive statistics for each activity category for the sample.

Number and type of activities reported vary throughout the day, as shown in Figure 5.2a. Most activities are reported in the early evening (5–9 p.m.), followed by the morning period (5–9 a.m.). For electricity research this is helpful since early evenings also have the highest electricity use. Direct comparison of the early evening period, which is recorded on the first and second day, shows that the number of activities reported reduces by around

Description	Response	Sample frequency [%]	Population frequency [%]
Home type	**Flat/apartment**	20	21
	Detached	15	17
	Semi-detached	23	25
	Terraced and other	42	37
Tenure	**Rent**	20	32
	Own	80	64
No. of rooms	2 or fewer	6	4
	3–5	59	54
	6 or more	35	42
Household income	<£25,000	18	41
	<£35,000	8	16
	<£50,000	18	17
	>£50,000	56	27
No. of occupants	1–2	55	64
	3–4	41	30
	>4	4	7
Age	Under 18	30	23
	19–50	52	44
	Over 50	18	35
Pet ownership	Cat	24	
	Dog	9	
	Fish	5	
Estimated monthly electric bill	<£400	40	
	<£700	41	
	>£700	19	
Bill affordability	Not affordable	15	
	Somewhat affordable	52	
	Affordable	33	
Electricity tariff	Renewable electricity	18	
	Economy 7 or 10	8	
	Other tariff	14	
	Do not know	14	
Large appliances	Washing machine	83	83
	Tumble dryer	31	50
	Washer dryer	16	14
	Underfloor heating	19	
	Gas boiler	83	85
	Heat pump	2	
	Electric stovetop	34	38
	Electric vehicle	0.3	0.4
	Electricity display	9	
	Solar thermal	8	
Other appliances (count)	TV/computer screens	M=2.39, SD=1.33	
	Dehumidifiers	M=0.17, SD=0.42	
	Air conditioners	M=0.06, SD=0.28	
	Portable heaters	M=0.33, SD=0.60	
	Night storage heaters	M=0.12, SD=0.67	
	Power showers	M=0.32, SD=0.65	
Season participated	**Winter (Nov-Feb)**	46	
	Shoulder (Mar–Apr; Sep–Oct)	19	
	Summer (May–Aug)	35	
Day of week participated	Sunday–Monday	17	
	Monday–Tuesday	16	
	Tuesday–Wednesday	8	
	Wednesday–Thursday	36	
	Thursday–Friday	8	
	Friday–Saturday	16	

Table 5.1: Selected socio-demographic, dwelling, appliance, and study participation date characteristics. National figures are based on estimates from various sources (CCC, 2016; DECC, 2013; ONS, 2016, 2017, 2018*a*,*b*).

Time-use category	Example activities	Mean (SD)
Care for home	Arrange things, clear up, washing machine, vacuum cleaning, using tumble dryer	3.0 (2.7)
Care for others	Put child to bed, play with child, being with others, bathing children, with pet	3.2 (3.4)
Care for self	Sleep, shower, getting dressed, in bed, getting ready	7.7 (5.8)
Food	Hot drink, eat hot meal, cook on hob, wash dishes, eating	12.5 (8.2)
Recreation	Watching TV, reading, socialising, Internet, E-mail	7.3 (6.2)
Work	Work, main job, work second job, work with machinery, study	0.6 (1.1)

Table 5.2: Activity categories, examples of their most frequently reported activities, and descriptive statistics.

23%. Respondent fatigue is a possible explanation for some participants. The relative share of activities reported in each category remains similar, however, so we do not expect this to bias the second 5–9 p.m. model.

Following Mckenna et al.'s (2017) point that the activity taxonomy developed for HETUS does not differentiate between 'energy-intensive' and 'low-energy' activities, we address this limitation by repeating the analysis described below including only activities we designate as 'energy-intensive'. Table 5.3 presents a list of these 26 activities and their descriptive statistics. These activities were reported a total of 2,218 times by study participants. In Figure 5.3, we show how the frequency of these activities varies throughout the day. While the total number of activities reported during each four-hour period remains proportional to the relative totals in Figure 5.2a, the share of 'food' and 'recreation' activities is larger, and no 'work' or 'other care' activities are reported.

5.3.5 Dependent variable: Average electricity demand and average 'de-minned' demand

We use several dependent variables in the following analyses. First, we model average daily electricity demand in Watts (W) across the full 28-hour study period. Next, we

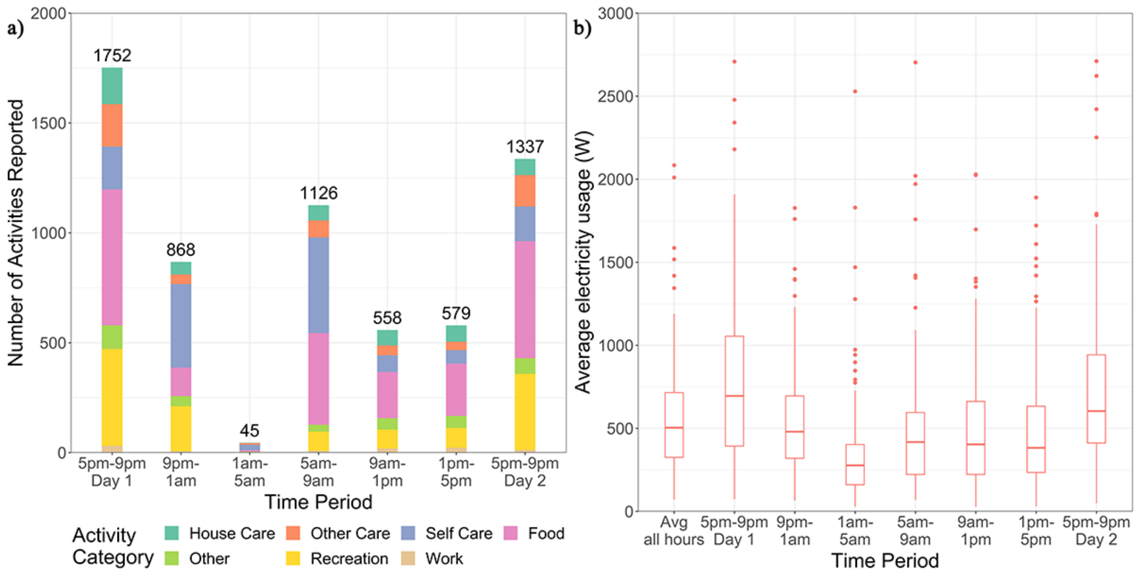

Figure 5.2: **a)** Frequency of activities reported by time period and time-use category ($N = 6265$). **b)** Boxplots of the sample's average electricity demand by time period ($N = 173$).

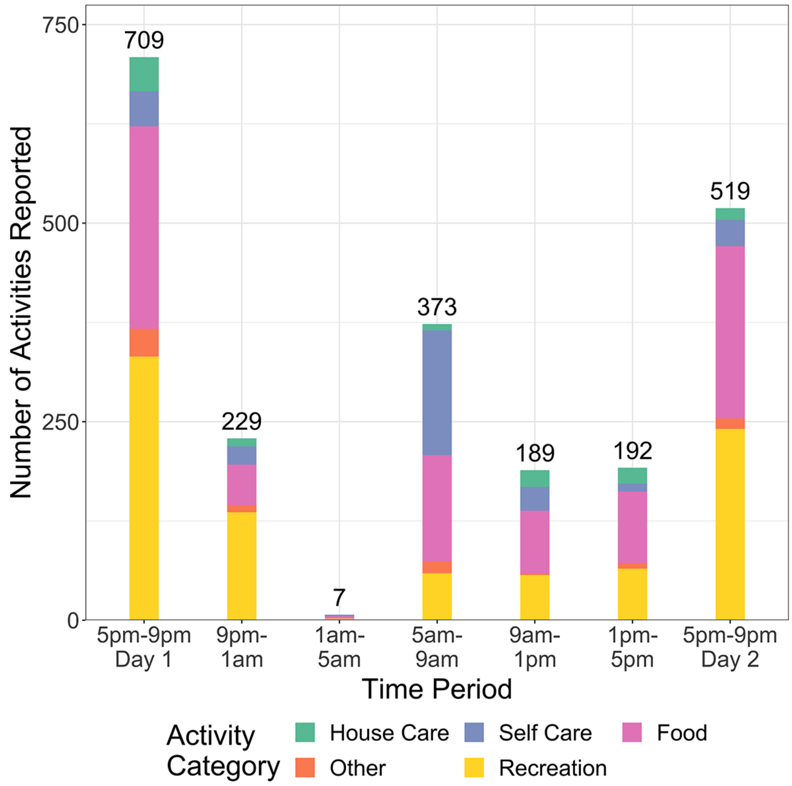

Figure 5.3: Frequency of 'energy-intensive' activities reported by time period and time-use category ($N = 2218$)

Time-use category	Time-use code	Activity description	Category mean (SD)
Care for home	3212	Vacuum cleaning	0.67 (1.02)
	3311	Washing machine	
	3320	Ironing	
Care for self	310	Wash and dress	1.73 (1.57)
	311	Shower	
	315	Hygiene/beauty	
	316	Wash	
	321	Bathing	
	8574	Shave	
Food	212	Eating hot meal	4.79 (3.70)
	214	Hot drink	
	3110	Prepare meal	
	3112	Prepare hot meal	
	3120	Baking	
	3124	Toaster	
	3133	Dish washer	
	3190	Prepare meal	
Recreation	3720	Internet	5.16 (4.63)
	3725	Online shopping	
	7170	E-mail	
	8001	Social media	
	8215	Watching TV	
	8221	Screen time	
	8552	Computer games	
	8553	TV	

Table 5.3: Time-use codes, categories, and activity descriptions for 'energy-intensive' activities.

model average daily 'de-minned' electricity demand for both the full study period and for four-hour periods throughout the day. Electricity readings are averaged for the time period to which each model corresponds (i.e., the full day models use the average and de-minned average from 5 p.m. on day 1 until 9 p.m. on day 2 while the other models use de-minned average electricity demand over each specified four-hour period). De-minning subtracts each household's minimum demand from its average demand in order to remove its 'baseload' (Jin et al., 2017). We use this technique because our aim is to characterise intra-day variations in electricity demand and especially those variations that can be explained by activity data, which we expect to show stronger associations to de-minned electricity demand than normal averages.

All dependent variables are log-transformed prior to regression analysis in order to improve linearity of regression residuals. Figure 5.2b shows boxplots for each dependent variable prior to de-minning and transformation. Boxes indicate 25th, 50th, and 75th percentiles, while the whiskers indicate 1.5 times the interquartile range below or above the 25th or 75th percentiles, respectively. Dots represent outliers beyond these. Mean demand during peak times is characteristically high compared to other day times. Some of the outliers in Figure 5.2b are households with unusually high demand throughout the day. As Figure 5.2 suggests, number of activities reported and average electricity demand are positively correlated during peak times, implying that activity reporting itself may be a useful proxy for understanding high demand.

The sample's mean daily electricity average of 565 W ($SD = 340W$) is slightly higher than the 'medium' estimates given by Ofgem's Typical Domestic Consumption Values (TDCV) for different classes of residential customers, which range from 350 to 490 W (Ofgem, 2017). This may be because most households take part in winter months when electricity use is typically higher and because of the sample biases previously mentioned.

5.3.6 Statistical analysis: Lasso and OLS regression

We use a variant on multiple linear regression to select final models of electricity demand. This variant is the 'least absolute shrinkage and selection operator' (lasso), developed by Tibshirani (1996). Lasso is one of several prominent regularisation techniques that have gained popularity for use in regression analyses where one or more of the following situations is present: the set of possible predictors is large and the analysis aims to identify those that most contribute to variations in the response variable; the data are 'high-dimensional', meaning the number of predictors is larger than the number of observations; or the data suffer from a high degree of multicollinearity between predictors, which occurs when one predictor variable can be linearly predicted from the others in the model.

These situations are often present in electricity demand modelling. Failing to address them can lead to poor performance of the model and biased results due to unreliable model coefficients. Furthermore, conventional methods for variable selection, such as the stepwise methods, have been shown to cause various analytical issues, including overestimated R-squared values and regression coefficients and predicted values that are too narrow (Harrell, 2001). Regularisation techniques address these issues by constraining the magnitudes of coefficients and by trading a small increase in model bias for a larger reduction in variance (see Satre-Meloy (2019) for more details). While regularisation techniques are the subject of a rich literature in statistics and machine learning and are well-suited to challenges in statistical modelling of electricity demand, their application in the energy literature thus far has been limited (Hsu, 2015; Huebner et al., 2016).

In this study, the data are neither high-dimensional ($n \gg p$) nor do the predictors suffer from a high degree of multicollinearity, which is determined by inspecting the variance inflation factors (VIFs) for the data.[1] For these reasons, the primary interest in applying regularisation is to perform variable selection for a large predictor set to achieve a sparser model. While one of the other fundamental regularisation techniques, ridge regression, performs well when dealing with high-dimensional data, it does not perform variable selection. The elastic net, a combination of ridge and lasso approaches, can address issues of multicollinearity but is somewhat more complex to implement (Zou and Hastie, 2005). Our motivation in applying lasso regression is thus to attain sparsity and deliver statistical and computational gains (Tibshirani, 1996).

Lasso applies a penalised linear regression model that shrinks the coefficients of some regression covariates while setting others exactly to zero, thus performing variable selection. While OLS regression estimates coefficients by minimising the residual sum of squares (RSS), lasso minimises the RSS with an added penalty parameter based on $\sum_{j=1}^{p} \mid \beta_j \mid$ for some multiplier λ. Minimising the RSS is thus given by the equation:

[1] VIFs signal whether regression coefficients are inflated due to correlation between predictor variables. If they are uncorrelated, VIF=1. Traditionally, VIFs greater than 10 indicate high multicollinearity (Roberts and Thatcher, 2009). In the models presented here, no VIFs are found to be greater than 10.

$$RSS(\beta) = \sum_{i=1}^{n} \left(y_i - \beta_0 - \sum_{j=1}^{p} \beta_j x_{ij} \right)^2 + \lambda \sum_{j=1}^{p} \mid \beta_j \mid \tag{5.1}$$

The amount of penalty or shrinkage that is applied to the regression coefficients is controlled by the parameter λ. As λ is increased, a larger penalty is applied, and the estimates are progressively shrunk toward zero. In this way, lasso enables the selection of a model that does not overfit the data but still has low error. Both the predictors and the response are standardised to have mean zero and a standard deviation of one prior to running lasso.

Estimating the optimal value for λ is done using k-fold cross-validation, where models are constructed using a range of values for λ and each model's mean-squared error (MSE) after cross-validation is plotted. Comparing the MSE for each model at varying λ values then enables the selection of a parsimonious model with low MSE, a process known as hyperparameter tuning in machine learning disciplines. In cases where model sparsity is a primary aim, it is common to follow the 'one standard error' rule, which says to select the simplest model that has an MSE within one standard error of the model with the minimum MSE (Friedman et al., 2010, p. 17). We follow this rule in the subsequent analyses to improve model interpretability.

Lasso regression is applied to the full predictor set in order to construct a model for full-day average electricity demand. The predictor variables selected at this stage are then included in an OLS regression, where both unstandardised and standardised coefficients are given. Unstandardised coefficients are measured in the original units of each independent variable. When the dependent variable has been log-transformed, it changes by 100 x (unstandardised coefficient) percent on average for each one unit increase in the predictor variable, holding all other variables in the model constant. Standardised coefficients are instead measured in terms of standard deviation and thus can be compared in magnitude to determine which predictors have stronger or weaker associations with demand given that predictors are measured in different units and on varying scales. Because our data contain numerous binary or dummy-coded predictors, we follow Gelman's (2008) suggestion of standardising all non-binary predictors by dividing by two standard deviations rather than one. This allows for numeric predictors' coefficients to be interpreted on the same scale as the binary inputs, which we leave unstandardised.

Next, we model daily average de-minned electricity demand, again using lasso to select influential predictors. Finally, we use lasso to select separate models for each four-hour time period, starting with 5–9 p.m. on day 1 and ending with 5–9 p.m. on day 2. We choose a four-hour duration for each model to balance the aim of modelling typical patterns of daily life (i.e., peak demand hours, late evening, overnight, early morning, afternoon, etc) while keeping the analysis from becoming overly complex, but our approach could be applied to averages taken over shorter or longer time durations. We investigate the predictors selected in each model and compare their coefficients in order to identify which predictors are influential in explaining demand patterns during different times of day. Next, using only activities we designate as 'energy-intensive', we repeat this analysis to examine whether using a more energy-relevant categorisation of activities improves model coefficients and/or explanatory power.

Finally, to investigate how much activity data contribute to explaining variations in de-minned average electricity demand (for both 'all' as well as 'energy-intensive' activities), we compare adjusted R-squared values for models with and without the activity variables included. *R Statistics* (R Core Team, 2017) and the associated packages *dplyr* (Wickham et al., 2019) and *glmnet* (Friedman et al., 2010) are used for data cleaning and regularised regression analyses. *Arm* (Gelman, 2008) is used to compute standardised coefficients, and *ggplot2* (Wickham, 2016) is used for plots.

5.4 Results

The following section presents regression results of full-day and four-hour models of average electricity demand using household survey and activity data.

Lasso regression is performed on all survey and activity data to select influential variables in models of daily average electricity demand and de-minned average electricity demand. The penalty parameter λ is tuned using 10-fold cross-validation. Figure 5.4 shows the cross-validation procedure. The model with the minimum MSE (left dotted line) is found with a λ value of 0.03, includes 23 predictors, and has an MSE of 0.25. The model with an MSE within one standard error of the minimum (right dotted line) is found with a λ value of 0.08 and includes 13 predictors.

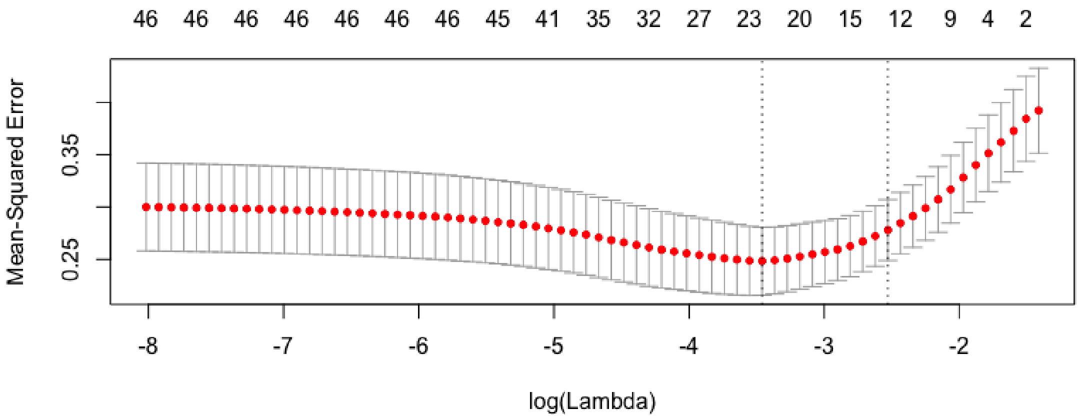

Figure 5.4: Plot of cross-validation MSE for lasso regression of daily average electricity demand. The horizontal bottom axis shows the logarithm of the tuning parameter λ, which increases in magnitude from left to right. The top horizontal axis shows the number of non-zero coefficients for each model run, and points and error bars show the mean and standard error of the cross-validation MSE, respectively. The left vertical dotted line indicates the λ value at which the minimum MSE is found, while the right vertical dotted line indicates the model with the fewest non-zero coefficients within one standard deviation of the minimum.

First, we include these selected 13 predictors in an OLS regression. Second, we re-run lasso on the full data, this time using de-minned average daily electricity demand as the dependent variable. For the second model, lasso selects two new variables and sets two from the first model to zero, meaning the total number of selected predictors remains at 13. We again include these variables in a simple OLS regression. Table 5.4 presents regression results, including unstandardised coefficients (B) and 95% confidence intervals along with standardised coefficients (β) for both full-day models, with coefficients listed in order of standardised coefficient magnitude. For all regression models, we examine diagnostic plots of fitted values versus residuals and Q–Q plots. With few exceptions, these plots confirm normality and linearity of the residuals.

Model 1 in Table 5.4 explains $R^2 = 49\%$ (adjusted $R^2 = 0.44$) of the variance in average electricity demand, $F(13, 159) = 10, p < 0.001$. The strongest predictors of increased daily demand are mostly dwelling and appliance-related variables, especially EV ownership, number of power showers, living in a detached home, number of rooms, and number of TVs/computers. Socio-demographic variables that correlate to increased demand include cat ownership and number of occupants. The only activity category variable selected is recreation; number of recreation activities reported is associated with increased daily average demand.

Ownership of a gas boiler and being on a renewable or 'green' electricity tariff are the only variables in the model that predict decreases in average daily demand. The former effect is likely observed because not owning a gas boiler indicates higher use of electricity for space and water heating. The latter effect may reflect a predisposition to conserve electricity among those who opt in to a renewable tariff.

The second model in Table 5.4, for which the dependent variable is de-minned average daily electricity demand, explains $R^2 = 43\%$ (adjusted $R^2 = 0.38$) of the variance of the dependent variable, $F(13, 159) = 9.2, p < 0.001$. In addition to the variables in model 1, lasso selects being on an Economy 7 or 10 tariff, which typically charge lower prices for nighttime or 'off-peak' electricity use, and number of night storage heaters. Both have positive coefficients but high uncertainty. This model drops EV ownership and number of rooms as predictors.

In general, the differences between models 1 and 2 are slight, with model 2 showing slightly lower explanatory power than model 1. Notably, we do not find additional activity variables are selected in model 2, even though we might suspect activities to show a stronger association to de-minned rather than average daily demand. We return to this finding in

Section 5.5.

	1) Daily average (W)		2) De-minned daily average (W)	
Predictor	B (5–95% CI)	β	B (5–95% CI)	β
EV	0.61** (0.18, 1.04)	0.61		
Tariff: Renewable	−0.30** (−0.50, −0.11)	−0.30	−0.30* (−0.54, −0.06)	−0.26
Gas boiler	−0.27* (−0.48, −0.06)	−0.27	−0.19 (−0.46, 0.09)	−0.16
Owns cat	0.25** (0.08, 0.43)	0.25	0.18 (−0.04, 0.39)	0.24
Tariff: Economy 7 or 10			0.43 (−0.03, 0.88)	0.24
Recreation activities	0.02** (0.01, 0.03)	0.24	0.03*** (0.01, 0.04)	0.23
No. power showers	0.15* (0.03, 0.27)	0.19	0.13 (−0.02, 0.28)	0.18
Detached home	0.19 (−0.04, 0.43)	0.19	0.45** (0.18, 0.73)	0.34
No. rooms	0.07* (0.005, 0.14)	0.18		
No. TV/computers	0.06* (0.004, 0.12)	0.17	0.06 (−0.02, 0.13)	0.18
Under floor heating	0.16 (−0.03, 0.35)	0.16	0.16 (−0.07, 0.40)	0.20
Tumble dryer	0.10 (−0.07, 0.27)	0.10	0.10 (−0.11, 0.30)	0.08
No. of night storage heaters			0.07 (−0.11, 0.25)	0.01
Estimated monthly electric bill	0.03 (−0.02, 0.09)	0.10	0.02 (−0.05, 0.09)	0.13
No. of occupants	0.03 (−0.04, 0.10)	0.08	0.08 (−0.01, 0.16)	0.11
Constant	5.42*** (5.10, 5.74)	6.27	4.97*** (4.61, 5.33)	6.15
Observations	173		173	
R^2	0.49		0.43	
Adjusted R^2	0.44		0.38	
Residual std. error (df = 159)	0.47		0.57	

Note: *$p < 0.05$, **$p < 0.01$, ***$p < 0.001$.

Table 5.4: Regression results for average electricity demand during full 28-hour study period. Included variables for each model are those selected using lasso regression. Model coefficients and confidence intervals are those resulting from an OLS regression of the lasso-selected variables.

To further investigate how categories of time-use activities might explain patterns in household electricity demand, we model de-minned average electricity demand over four-hour time periods. Lasso regression is used to select models (which can have varying numbers of predictors), and we again include lasso-selected variables in an OLS regression for each four-hour period to investigate model coefficients. Table 5.5 shows variables and unstandardised coefficients with 95% confidence intervals. We exclude standardised coefficients from this table for simplicity but include these in Table 5.A.1 in the appendix.

Table 5.5 shows that for all times of day except 1–5 a.m., activity category variables are frequently selected in these models and show strong associations to de-minned electricity demand. Activities in the 'care for home', 'food', and 'recreation' categories, in particular, are consistently selected as influential predictors across four-hour models.

Coefficients for activity variables show how different types of activities vary in their associations to electricity demand at different times of day. Care for the home, which includes some energy-intensive household chores, has a stronger association with de-minned demand in the early evening, mid-morning, and afternoon models and a weaker association in the late evening and early morning models. It is absent from the overnight model (1–5 a.m.) as well as from the day 2 5–9 p.m. model. Food-related activities follow patterns of typical mealtimes, with stronger coefficients in the evenings, early morning, and afternoon. The models do not show a consistent result in terms of the influence of different mealtimes on demand, as the food coefficient for the morning mealtime is larger than in the day 1 5–9 p.m. model but smaller than in the day 2 5–9 p.m. model. This finding highlights that electricity demand of food preparation and consumption can vary between days during the same mealtimes. Recreation activities show strong associations to electricity demand except overnight and during the early morning, but this variable is also absent from the day 2 5–9 p.m. model.

The only model in which no activity variables are selected is the overnight model, which also happens to be the model with the fewest selected predictors. Only EV ownership and being on an Economy 7 or 10 tariff are selected. Both of these can reasonably be expected to explain increases in nighttime electricity demand.

Regarding non-activity predictors that are selected in the four-hour models, many of these variables are the same as those selected in the models of 28-hour electricity demand. Some of these variables, such as living in a detached home, number of occupants, and estimated monthly electric bill appear to influence de-minned electricity consumption consistently throughout the day, as indicated by their selection in the majority of the models. Other variables are much more time-specific in their associations with de-minned electricity demand. For instance, gas boiler ownership shows a large, negative association

Predictor	5–9 p.m. (1)	5–9 p.m. (2)	9 p.m.–1 a.m.	1 a.m.–5 a.m.	5 a.m.–9 a.m.	9 a.m.–1 p.m.	1 p.m.–5 p.m.
Activity categories							
Care for home	0.17*** (0.09, 0.24)		0.16* (0.002, 0.32)		0.16 (−0.07, 0.39)	0.15 (−0.02, 0.33)	0.15 (−0.03, 0.32)
Care for others					0.15 (−0.04, 0.33)	0.35** (0.10, 0.60)	
Food	0.03 (−0.02, 0.07)	0.10*** (0.05, 0.15)	0.08* (0.01, 0.16)		0.06 (−0.01, 0.13)		0.08 (−0.02, 0.19)
Other activity							0.22* (0.02, 0.43)
Recreation	0.06* (0.01, 0.12)		0.08* (0.01, 0.15)			0.21* (0.02, 0.39)	0.25* (0.04, 0.46)
Appliances							
EV				2.23*** (1.06, 3.40)			
Gas boiler			−0.37* (−0.65, −0.09)		−0.59** (−0.99, −0.20)		
Heat pump					0.76 (−0.20, 1.72)		
No. night storage heaters							0.39** (0.15, 0.63)
No. power showers			0.23** (0.06, 0.39)			0.27* (0.01, 0.53)	
No. TVs/computers		0.07 (−0.04, 0.18)				0.17** (0.04, 0.30)	0.09 (−0.04, 0.22)
Tumble dryer					0.28 (−0.04, 0.60)		0.37* (0.01, 0.74)
Under floor heating			0.11 (−0.15, 0.37)		0.34 (−0.04, 0.72)		
Dwelling							
Detached home	0.62*** (0.29, 0.95)	0.56** (0.15, 0.97)	0.20 (−0.10, 0.50)		0.17 (−0.25, 0.59)		0.66** (0.19, 1.14)
No. rooms in home		0.08 (−0.05, 0.20)				0.10 (−0.04, 0.23)	0.08 (−0.06, 0.22)
Terraced home	−0.16 (−0.41, 0.08)						
Socio-demographics							
Bill affordability	−0.18* (−0.35, −0.01)						
Estimated monthly electric bill		0.09 (−0.01, 0.19)	0.06 (−0.02, 0.13)		0.09 (−0.02, 0.19)	0.08 (−0.04, 0.21)	
No. of occupants	0.13* (0.03, 0.22)		0.06 (−0.03, 0.15)		0.09 (−0.06, 0.23)		0.13 (−0.02, 0.27)
Owns cat		0.31 (−0.01, 0.64)	0.36** (0.12, 0.59)		0.35* (0.02, 0.69)		
Participated in summer	−0.43*** (−0.67, −0.18)						
Participated in spring/autumn			0.34** (0.09, 0.59)				
Tariff: Economy 7 or 10				1.73*** (1.07, 2.40)			
Tariff: Renewable	−0.42** (−0.72, −0.11)		−0.43** (−0.69, −0.17)				
Tenure (rent vs. own)					−0.38* (−0.73, −0.02)		
Constant	5.82 (5.38, 6.26)	4.55 (4.00, 5.10)	5.04 (4.68, 5.40)	3.92 (3.75, 4.09)	4.53 (3.98, 5.07)	3.53 (2.91, 4.15)	3.42 (2.79, 4.04)
Observations	171	170	172	172	172	170	170
R^2	0.43	0.25	0.43	0.28	0.33	0.29	0.39
Adj R^2	0.40	0.23	0.39	0.27	0.28	0.26	0.35

Note: $^{*}p < 0.05$, $^{**}p < 0.01$, $^{***}p < 0.001$.

Table 5.5: Regression results for de-minned average electricity demand: four-hour models (5 p.m. day 1 until 9 p.m. day 2). Included variables for each model are those selected using lasso regression. Model coefficients and confidence intervals are those resulting from an OLS regression of the lasso-selected variables.

with electricity demand in the late evening and early morning models. This finding may be intuitive, since space heating is typical in the late evening and early morning, and homes that heat with gas rather than electricity will likely have lower electricity demand during these hours.

Because households select their own date for participation, we include season and day of week as candidate predictors for each model. Table 5.5 shows that there does not appear to be a notable day-of-week effect for our sample, but we do find a slight seasonal effect, as the 5–9 p.m. day 1 model includes the ‘participated in summer’ predictor with a negative coefficient, and the 9 p.m.–1 a.m. model includes the ‘participated in spring or autumn’ predictor with a positive coefficient. These coefficients can be interpreted in terms of their comparison to the reference category, which is winter participation. Households that participate in summer show lower 5–9 p.m. de-minned electricity demand than households that participate in winter; similarly, households that participate in spring/autumn show slightly increased 9 p.m.–1 a.m. demand compared with households participating in winter. Regarding model explanatory power, we find that the model with the largest proportion of explained variance is the 5–9 p.m. day 1 model ($R^2 = 43\%$; adjusted $R^2 = 0.40$) and that the model with the lowest is the 5–9 p.m. day 2 model ($R^2 = 25\%$; adjusted $R^2 = 0.23$). The results for the two comparable 5–9 p.m. models also show inconsistencies in terms of variables selected, especially in the case of the activity variables. This finding challenges the notion of ‘average load profiles’ (as commonly used for settlement, see Elexon (2013)) and suggests that activity-related factors that influence electricity demand vary from day to day. This is a preliminary finding, however, and draws from a limited sample of days.

To address the limitation that these categories of activities do not differentiate between ‘energy-intensive’ and ‘low-energy’ activities, we repeat this analysis including only those activities that we designate as ‘energy-intensive’. Table 5.6 presents results for these models (with Table 5.A.2 in the appendix displaying standardised coefficients). With the exception of several coefficients, we observe in most cases a marked increase in the care for home, food, and recreation activity variables’ coefficients. This is especially true for ‘care for home' activities reported in the morning, mid-morning, and afternoon periods, suggesting energy-intensive housework is a strong predictor of de-minned demand during these times.

In some instances, the energy-intensive activity models either drop or add activity-related variables for different times of day. While the 'care for home' category when not filtering for energy-intensive activities is selected in the day 1 5–9 p.m. and the 9 p.m.–1 a.m. models, it is not selected when only energy-intensive activities are included. This may suggest that more energy-intensive activities related to household chores are not as common in the evenings. It also indicates that 'care for home' as a category includes non-energy-intensive activities that are important for explaining electricity demand during evenings.

Similar distinctions are found for food-related activities. The model including only energy-intensive food activities more clearly differentiates between evening and morning meals, as morning food activities show less associations with electricity demand when excluding non-energy-intensive activities, likely because morning meals are often cold and do not involve appliance use. For recreation activities, the timing of when these are influential does not change between the two different categorisations, though the strength of associations clearly increases when only energy-intensive activities are included.

To test the extent to which categorised activity data improve model explanatory power, we compute the adjusted R-squared value for each model excluding the lasso-selected activity predictors for that model, and we do this for both categorisations tested. We present these results in Table 5.7. We see that excluding activity data reduces model explanatory power by an average of nine percentage points across models (excluding the 1–5 a.m. model). We do not observe, however, a noticeable difference between the models including all activity data and those including only 'energy-intensive' activities. These findings provide evidence of the potential improvements to household electricity modelling by including activity data, even when it is relatively coarsely categorised.

Predictor	5–9 p.m. (1)	5–9 p.m. (2)	9 p.m.–1 a.m.	1 a.m.–5 a.m.	5 a.m.–9 a.m.	9 a.m.–1 p.m.	1 p.m.–5 p.m.
Activity categories							
Care for home					1.05** (0.39, 1.71)	0.52* (0.13, 0.91)	0.70** (0.24, 1.15)
Care for self					0.11 (−0.05, 0.27)		
Food	0.02 (−0.06, 0.11)	0.18** (0.06, 0.30)	0.23** (0.06, 0.40)			0.14 (−0.07, 0.35)	0.20* (0.01, 0.39)
Recreation	0.10** (0.04, 0.16)		0.20*** (0.10, 0.29)			0.35** (0.12, 0.58)	0.30* (0.05, 0.55)
Appliances							
EV				2.23*** (1.06, 3.40)			
Gas boiler			−0.32* (−0.59, −0.05)		−0.65*** (−1.03, −0.27)		
No. night storage heaters							0.42*** (0.18, 0.66)
No. power showers			0.26** (0.10, 0.42)			0.22 (−0.05, 0.49)	
No. TVs/computers						0.14* (0.01, 0.27)	0.08 (−0.05, 0.20)
Tumble dryer						0.34 (−0.03, 0.71)	0.38* (0.02, 0.75)
Under floor heating			0.12 (−0.13, 0.38)		0.40* (0.03, 0.77)	0.38 (−0.05, 0.81)	
Dwelling							
Detached home	0.68*** (0.34, 1.03)	0.66** (0.24, 1.07)	0.20 (−0.09, 0.50)		0.28 (−0.14, 0.69)	0.18 (−0.32, 0.69)	0.76** (0.28, 1.24)
No. rooms in home		0.08 (−0.04, 0.21)				−0.01 (−0.16, 0.14)	0.07 (−0.07, 0.21)
Terraced home	−0.13 (−0.39, 0.13)						
Socio-demographics							
Estimated monthly electric bill		0.11* (0.01, 0.21)	0.04 (−0.03, 0.11)		0.10 (−0.01, 0.20)	0.04 (−0.08, 0.17)	
No. of occupants	0.16** (0.06, 0.26)		0.08 (−0.01, 0.17)		0.15* (0.01, 0.29)	0.12 (−0.04, 0.28)	0.15* (0.001, 0.30)
Owns cat			0.33** (0.10, 0.56)		0.33 (−0.01, 0.67)		
Participated in summer	−0.43** (−0.69, −0.17)						
Participated in spring/autumn			0.36** (0.12, 0.61)				
Participated on Fri-Sat						0.56* (0.11, 1.01)	
Tariff: Economy 7 or 10				1.73*** (1.07, 2.40)			
Tariff: Renewable	−0.51** (−0.81, −0.20)		−0.40** (−0.65, −0.14)				
Constant	5.56(5.24, 5.87)	4.77 (4.24, 5.31)	4.99 (4.63, 5.34)	3.92 (3.75, 4.09)	4.52 (4.03, 5.00)	3.62 (2.97, 4.28)	3.48 (2.86, 4.11)
Observations	171	170	172	172	172	170	170
R^2	0.35	0.20	0.45	0.28	0.30	0.33	0.37
Adj R^2	0.32	0.18	0.41	0.27	0.26	0.28	0.34

Note: $*p < 0.05$, $**p < 0.01$, $***p < 0.001$.

Table 5.6: Regression results for de-minned average electricity use: four-hour models (5 p.m. day 1 until 9 p.m. day 2) including only 'energy-intensive' activities across categories. Included variables for each model are those selected using lasso regression. Model coefficients and confidence intervals are those resulting from an OLS regression of the lasso-selected variables.

	5–9 p.m. (1)	5–9 p.m.(2)	9 p.m.–1 a.m.	1 a.m.–5 a.m.	5 a.m.–9 a.m.	9 a.m.–1 p.m.	1 p.m.–5 p.m.
'All' activity models							
Adj R^2	0.40	0.23	0.39	0.27	0.28	0.26	0.35
Adj R^2 without activity variables	0.28	0.17	0.30	–	0.25	0.13	0.23
Difference in adj R^2	0.12	0.06	0.09	–	0.03	0.13	0.12
'Energy-intensive' activity models							
Adj R^2	0.31	0.18	0.41	0.27	0.26	0.28	0.34
Adj R^2 without activity variables	0.27	0.15	0.30	–	0.22	0.18	0.23
Difference in adj R^2	0.04	0.03	0.11	–	0.04	0.10	0.11

Table 5.7: Model explanatory power with and without lasso-selected activity variables for both 'all' activity data and 'energy-intensive' activity data models.

5.5 Discussion

5.5.1 Summary and discussion of results

The results show that 13 lasso-selected predictors can explain 49% of full-day average household electricity demand. When the dependent variable is de-minned to remove baseload demand, lasso selects a model that explains 43% of the variance in average de-minned electricity demand. These results are consistent with or in some cases improve on model explanatory power found in previous studies of household electricity usage (Huebner et al., 2016; Kavousian et al., 2013; McLoughlin et al., 2012). For the models of four-hour average de-minned electricity demand, explanatory power is somewhat lower, ranging from 25–43% of variance explained. These models include varying numbers of lasso-selected predictors, but we find that removing activity variables from these models reduces their explanatory power considerably.

The results for both full-day models confirm previous findings that appliance ownership and occupant socio-demographics show the strongest associations to electricity demand patterns (Huebner et al., 2016). These results also, however, identify new variables that might help predict daily demand, such as electricity tariff choice or pet ownership. The results provide some evidence that these may be important for understanding electricity demand patterns.

This study goes beyond previous research by investigating how the associations between appliance, dwelling, and socio-demographic variables and de-minned electricity demand vary as a function of time of day. The EV ownership variable, for instance, is only selected in the nighttime model, which suggests that ownership of an EV increases a household's nighttime demand more so than it increases demand during other times. Similarly, the number of power showers and presence of underfloor heating are associated with demand during late evening and morning hours, when showers are likely more frequent. Number of TVs/computers owned is a stronger predictor in the morning and evening.

The study also tests the extent to which different types of activities are associated with electricity demand during different periods of the day. We find evidence that certain types of activities are more relevant for understanding electricity demand patterns and also that certain activities have stronger or weaker associations at different times of day.

In particular, our results show that housework, eating and meal preparation, and recreation or media use are strongly associated with de-minned electricity demand, confirming similar findings from previous research (Anderson and Torriti, 2018; De Lauretis et al., 2017; Jalas and Juntunen, 2015; Palmer et al., 2013). The times at which these activities are stronger predictors of electricity demand also provides insight into when activities and electricity use are more tightly coupled, such as during evening peak periods and later evenings for food and recreational activities, early mornings for personal care activities, and early and mid-mornings for household chores.

Our results also highlight the extent to which denoting activities as 'energy-intensive' can strengthen these links. We find in general a sizeable increase in the strength of relationships between number of activities reported and average de-minned electricity demand throughout the day when including only 'energy-intensive' activities. We do note, however, that these results are somewhat mixed. For instance, the 'care for home' variable is significant in the day 1 5–9 p.m. model when including all activity data but not when including only energy-intensive activities. This finding suggests other 'care for home' activities are important for explaining electricity demand during this time. While acknowledging that our designation of activities as 'energy-intensive' is subjective, we take this finding as evidence that the links between activities and electricity demand may be more complex than simply differentiating activities on the basis of their expected use of energy. Further discussion of this point can be found in Grunewald and Diakonova (2018*a*).

Two additional results merit further discussion. The first is our finding that we do not observe an increase in the number of activity variables selected between the daily average and daily de-minned average models. We do, however, observe frequent selection of activity variables in the four-hour de-minned demand models. We expect that this observation results from the duration of time over which we are de-minning electricity demand. Over a 28-hour period, de-minning does not appear to yield stronger associations between activities reported and average demand. Over shorter four-hour periods, however, de-minning does appear to strengthen relationships between reported activities and demand. This finding suggests that activity data can improve models of shorter duration more so than it can models of longer duration. We expect activities to show increasingly strong associations to more finely resolved electricity readings, and this is something we are exploring in current research.

The second result deserving of further discussion is our finding that the models for the day 1 5–9 p.m. period and the day 2 5–9 p.m. period show considerable differences in the variables selected and in model explanatory power. We note that fewer activities are reported during the day 2 5–9 p.m. period, but given that lasso standardises predictors prior to selection, this should not influence the algorithm's selection procedure. Furthermore, we note that the 1–5 p.m. model includes several activity variables and higher explanatory power than the day 2 5–9 p.m. model, even with far fewer activities reported. We take this result to suggest that the relationship between activities and electricity demand may vary considerably in the same households between days and also that peak hours may be especially variable. This finding, too, warrants further exploration in future research.

5.5.2 Policy implications

These results have implications both for energy demand models and for policy considerations surrounding demand flexibility and DSR. First, we have shown that activity data, whether categorised as 'energy-intensive' or not, can improve models of household electricity use. For at least one of our 'peak-period' models, we find that activity data are especially useful for understanding variations in de-minned electricity use during these hours.

Our novel approach to collect time-use data in parallel with electricity demand can overcome many of the limitations and challenges previously discussed in the literature on time-use or activity-based models of energy demand (McKenna et al., 2017; Torriti, 2014). Specifically, this method enables the collection of more energy-relevant activities, which we have shown to increase the strength and statistical validity of the relationships between household activities and electricity demand. Given the lack of empirical evidence on these relationships, these results can help improve the specification of activity-appliance signals in models incorporating time-use data. Our approach also facilitates the collection of activity data from multi-occupant households, which have previously been more difficult to account for in time-use models of electricity demand.

Second, while some literature has concluded that it is sufficient to model active occupancy states for the purpose of constructing more accurate energy demand models, we believe that such an approach fails to more directly link the types of activities that are being performed during 'active occupancy', which have important consequences and policy implications for delivering more flexible demand.

Our results show that certain types of activities have stronger relationships to electricity demand at different times of day. Targeting these activities in DSR interventions could yield larger shifts in demand. Furthermore, tailoring DSR interventions based on better evidence about what is actually occurring in households during peak times and about how these activities vary among different segments of the population can improve their effectiveness while also mitigating their impact on vulnerable populations.

In terms of potential for shifting demand alone, these results suggest food preparation and meal times, household chores, and recreational activities should be prioritised for activity-led DSR given their strong associations with electricity use. But this is not the only consideration upon which effective DSR strategies should be based. Other considerations, such as those comprising Toritti et al.'s (2015) 'Flexibility Index', are also valuable for determining which activities at which times can be shifted without adverse social impacts.

5.5.3 Study limitations

Some limitations are present in this analysis. The sample is non-representative, and several sample biases are present. The sample size is also small relative to the number of predictors. Standard errors for most variables in the regression models are therefore large in comparison to coefficient size. Models are not constructed for different days of week or seasons in order to preserve the sample size, but these variables are instead included at the model selection stage. Furthermore, self-reported activities can lead to biases and inaccuracies. Measuring activities as simple frequency counts is a coarse way to include these data and may obfuscate more complex relationships between activity type, timing, and electricity demand. A similar approach as the one taken by Rhodes et al. (2014) and McLoughlin et al. (2015), who cluster load profiles and use household characteristics to build predictive models of membership in distinctive load profiles, may be instructive in this context. Time-use data may be even better suited to this approach than socio-demographic, dwelling, or appliance ownership data, due to its temporal resolution.

Our initial categorisation of activity data relies on a taxonomy that was not developed for energy modelling purposes, and although we make an attempt at categorising activities as 'energy-intensive', this is a subjective exercise. We believe much more work could be done to identify more energy-relevant activity categories with useful implications for energy demand modelling and for estimating the potential of DSR (Anderson, 2016).

Finally, more data from a more representative sample of households could reduce uncertainty surrounding the strength of relationships between predictive variables and electricity demand and could make these results more generalisable. The data collection is ongoing, such that more detailed analyses can be conducted in future. We believe there is potential to scale this research to much larger ($> 1,000$ households), nationally-representative samples using app-based activity diaries coupled with smart meter data, and preparations are underway for scaling this study. We also encourage similar studies across varying cultural and societal contexts.

5.6 Conclusions

This paper presents analysis of intra-day electricity demand for a sample of 173 U.K. households. Electricity provision during peak times is costly and challenging in energy systems with high penetrations of renewable electricity. Understanding the drivers behind electricity demand during peak times can assist in the design of more effective intervention strategies to shift electricity use patterns in the residential sector.

We present regression models of average full-day and four-hour de-minned electricity demand using a wide range of predictive factors, including socio-demographic, physical dwelling, appliance ownership, and categorised activity variables. Our analysis tests the ability of these variables to explain demand at different times of day and the strength of these relationships.

Our results show that adding activity data to regression models can improve their explanatory power of de-minned electricity demand during most times of day. This result holds regardless of whether only activities that we reasonably assume to be 'energy-intensive' are included. Given how challenging it is to achieve greater explanatory power of highly diverse electricity uses in households, a nine-percentage point increase in adjusted R-squared on average is encouraging. The results also show strong associations between activities and electricity demand at different times of day. The importance of household chores, food, and recreation activities for explaining electricity demand patterns is clearly demonstrated.

We expect that more nuanced categorisations of activities and investigations of their relationship to electricity demand at more finely resolved timescales will yield improved insights into the flexibility of demand. This evidence could reduce the cost of integrating renewable energy sources into existing energy systems by providing a better estimate of the potential of DSR in the residential sector, thus contributing to accelerating the low-carbon energy transition.

Data availability

The data that support the findings of this study are available from the corresponding author upon request. The tools used for time-use and electricity data collection are publicly available at `https://goo.gl/81mvCB`.

Acknowledgements

The authors gratefully acknowledge funding for this research from the Engineering and Physical Sciences Research Council (EPSRC) under grant EP/M024652/1.

5.7 References

Anderson, B. (2016), 'Laundry, energy and time: Insights from 20 years of time-use diary data in the United Kingdom', *Energy Research & Social Science* **22**, 125–136.

Anderson, B., Lin, S., Newing, A., Bahaj, A. and James, P. (2016), 'Electricity consumption and household characteristics: Implications for census-taking in a smart metered future', *Computers, Environment and Urban Systems* .

Anderson, B. and Torriti, J. (2018), 'Explaining shifts in U.K. electricity demand using time use data from 1974 to 2014', *Energy Policy* **123**, 544–557.

BEIS (2017), The Clean Growth Strategy: Leading the way to a low carbon future, Technical report, Department for Business, Energy & Industrial Strategy, London, United Kingdom.

BEIS (2018), DUKES chapter 5: Statistics on electricity from generation through to sales, *in* 'Digest of UK Energy Statistics (DUKES): Electricity', Department for Business, Energy & Industrial Strategy, London, United Kingdom, pp. 111–153.

CCC (2016), Annex 2: Heat in U.K. buildings today, *in* 'Next Steps for UK Heat Policy', Committee on Climate Change, London, United Kingdom.

De Lauretis, S., Ghersi, F. and Cayla, J.-M. (2017), 'Energy consumption and activity patterns: An analysis extended to total time and energy use for French households', *Applied Energy* **206**, 634–648.

DECC (2012), Demand side response in the domestic sector - A literature review of major trials, Technical report, Department of Energy and Climate Change, London, United Kingdom.

DECC (2013), Report 9: Domestic appliances, cooking & cooling equipment, Technical Report BRE report number 288143, Department of Energy and Climate Change, London, United Kingdom.

DECC (2014), Seasonal variations in electricity demand, Technical report, Department of Energy and Climate Change, London, United Kingdom.

Druckman, A., Buck, I., Hayward, B. and Jackson, T. (2012), 'Time, gender and carbon: A study of the carbon implications of British adults' use of time', *Ecological Economics* **84**, 153–163.

Elexon (2013), Load profiles and their use in electricity settlement, Technical report, Elexon, London, United Kingdom.

Ellegård, K. and Palm, J. (2011), 'Visualizing energy consumption activities as a tool for making everyday life more sustainable', *Applied Energy* **88**(5), 1920–1926.

Ellegård, K., Vrotsou, K. and Widén, J. (2010), VISUAL-TimePAcTS/energy use – a software application for visualizing energy use from activities performed, *in* 'Proceedings of the 3rd International Scientific Conference on "Energy Systems with IT"', pp. 1–17.

eurostat (2009), *Harmonised European Time Use Surveys: 2008 Guidelines*, Office for Official Publications of the European Communities, Luxembourg. OCLC: 698893637.

Firth, S., Lomas, K., Wright, A. and Wall, R. (2008), 'Identifying trends in the use of domestic appliances from household electricity consumption measurements', *Energy and Buildings* **40**(5), 926–936.

Friedman, J., Hastie, T. and Tibshirani, R. (2010), 'Regularization paths for generalized linear models via coordinate descent', *Journal of Statistical Software* **33**(1), 1.

Gelman, A. (2008), 'Scaling regression inputs by dividing by two standard deviations', *Statistics in Medicine* **27**(15), 2865–2873.

Gershuny, J. (2008), Time-use studies: Daily life and social change: Full research report, ESRC End of Award Report, RES-000-23-0704-A, ESRC, Swindon, United Kingdom.

Gershuny, J. and Sullivan, O. (2017), 'United Kingdom Time Use Survey, 2014-2015. [data collection]'.

Grunewald, P. and Diakonova, M. (2018*a*), 'The electricity footprint of household activities–implications for demand models', *Energy and Buildings* **174**, 635–641.

Grunewald, P. and Diakonova, M. (2018*b*), 'Flexibility, dynamism and diversity in energy supply and demand: A critical review', *Energy Research & Social Science* **38**, 58–66.

Grünewald, P. and Diakonova, M. (2019), 'METER data', `http://www.energy-use.org/docs/data/`.

Grünewald, P., Diakonova, M., Zilli, D., Bernard, J. and Matousek, A. (2017), What we do matters – a time-use app to capture energy relevant activities, *in* 'Proceedings of the 2017 ECEEE Summer Study on Energy Efficiency in Buildings', European Council for an Energy Efficient Economy, Presqu'île de Giens, France, pp. 2085–93.

Grünewald, P. and Layberry, R. (2015), Measuring the relationship between time-use and electricity consumption, *in* 'Proceedings of the 2015 ECEEE Summer Study on Energy Efficiency in Buildings', European Council for an Energy Efficient Economy, Presqu'île de Giens, France, pp. 2087–96.

Harrell, F. (2001), *Regression Modeling Strategies: With Applications to Linear Models, Logistic Regression, and Survival Analysis*, Springer Series in Statistics, Springer-Verlag, New York, NY.

Hsu, D. (2015), 'Identifying key variables and interactions in statistical models of building energy consumption using regularization', *Energy* **83**, 144–155.

Huebner, G., Shipworth, D., Hamilton, I., Chalabi, Z. and Oreszczyn, T. (2016), 'Understanding electricity consumption: A comparative contribution of building factors, socio-demographics, appliances, behaviours and attitudes', *Applied Energy* **177**, 692–702.

Jalas, M. and Juntunen, J. K. (2015), 'Energy intensive lifestyles: Time use, the activity patterns of consumers, and related energy demands in Finland', *Ecological Economics* **113**, 51–59.

Jin, L., Lee, D., Sim, A., Borgeson, S., Wu, K., Spurlock, C. A. and Todd, A. (2017), Comparison of clustering techniques for residential energy behavior using smart meter data, *in* 'Proceedings of the Thirty-First AAAI Conference on Artifical Intelligence', Association for the Advancement of Artificial Intelligence, San Francisco, CA.

Jones, R. V., Fuertes, A. and Lomas, K. J. (2015), 'The socio-economic, dwelling and appliance related factors affecting electricity consumption in domestic buildings', *Renewable and Sustainable Energy Reviews* **43**, 901–917.

Jones, R. V. and Lomas, K. J. (2016), 'Determinants of high electrical energy demand in U.K. homes: Appliance ownership and use', *Energy and Buildings* **117**, 71–82.

Kavousian, A., Rajagopal, R. and Fischer, M. (2013), 'Determinants of residential electricity consumption: Using smart meter data to examine the effect of climate, building characteristics, appliance stock, and occupants' behavior', *Energy* **55**, 184–194.

Lader, D., Short, S. and Gershuny, J. (2006), The Time Use Survey, 2005: How we spend our time, Technical report, Office for National Statistics, London, United Kingdom.

McKenna, E., Higginson, S., Grunewald, P. and Darby, S. J. (2017), 'Simulating residential demand response: Improving socio-technical assumptions in activity-based models of energy demand', *Energy Efficiency* .

McKenna, E. and Thomson, M. (2016), 'High-resolution stochastic integrated thermal–electrical domestic demand model', *Applied Energy* **165**, 445–461.

McLoughlin, F., Duffy, A. and Conlon, M. (2012), 'Characterising domestic electricity consumption patterns by dwelling and occupant socio-economic variables: An Irish case study', *Energy and Buildings* **48**, 240–248.

McLoughlin, F., Duffy, A. and Conlon, M. (2015), 'A clustering approach to domestic electricity load profile characterisation using smart metering data', *Applied Energy* **141**, 190–199.

National Infrastructure Commission (2016), Smart Power, Technical report, National Infrastructure Commission, London, United Kingdom.

Ofgem (2010), Demand side reponse. A discussion paper, Technical report, Office of Gas and Electricity Markets, London, United Kingdom.

Ofgem (2017), Typical Domestic Consumption Values for gas and electricity, Technical report, Office of Gas and Electricity Markets, London, United Kingdom.

ONS (2016), 'Population estimates for U.K., England and Wales, Scotland and Northern Ireland: Mid-2015', `http://www.ons.gov.uk/peoplepopulationandcommunity/populationandmigration/populationestimates/bulletins/annualmidyearpopulationestimates/mid2015`.

ONS (2017), 'Families and households: 2017', `http://www.ons.gov.uk/peoplepopulationandcommunity/birthsdeathsandmarriages/families/bulletins/familiesandhouseholds/2017`.

ONS (2018*a*), English Housing Survey: Headline Report, 2016-17, Technical report, Ministry of Housing, Communities & Local Government, London, United Kingdom.

ONS (2018*b*), 'Percentage of households with durable goods, U.K.: Table A45', `http://www.ons.gov.uk/peoplepopulationandcommunity/personalandhouseholdfinances/expenditure/datasets/percentageofhouseholdswithdurablegoodsuktablea45`.

Palmer, J., Terry, N., Kane, T., Firth, S., Hughes, M., Pope, P., Young, J., Knight, D. and Godoy-Shimizu, D. (2013), Further analysis of the household electricity use survey - Electrical appliances at home: Tuning in to energy saving, Technical Report 475/09/2012, Cambridge Architectural Research Limited, elementenergy, and Loughborough University, London, United Kingdom.

Powells, G., Bulkeley, H., Bell, S. and Judson, E. (2014), 'Peak electricity demand and the flexibility of everyday life', *Geoforum* **55**, 43–52.

R Core Team (2017), 'R: A language and environment for statistical computing', R Foundation for Statistical Computing.

Rhodes, J. D., Cole, W. J., Upshaw, C. R., Edgar, T. F. and Webber, M. E. (2014), 'Clustering analysis of residential electricity demand profiles', *Applied Energy* **135**, 461–471.

Richardson, I., Thomson, M., Infield, D. and Clifford, C. (2010), 'Domestic electricity use: A high-resolution energy demand model', *Energy and Buildings* **42**(10), 1878–1887.

Roberts, N. and Thatcher, J. (2009), 'Conceptualizing and testing formative constructs: Tutorial and annotated example', *ACM SIGMIS Database* **40**(3), 9–39.

Satre-Meloy, A. (2019), 'Investigating structural and occupant drivers of annual residential electricity consumption using regularization in regression models', *Energy* **174**, 148–168.

Sekar, A., Williams, E. and Chen, R. (2018), 'Changes in time use and their effect on energy consumption in the United States', *Joule* **2**, 1–16.

Stankovic, L., Stankovic, V., Liao, J. and Wilson, C. (2016), 'Measuring the energy intensity of domestic activities from smart meter data', *Applied Energy* **183**, 1565–1580.

Strbac, G., Aunedi, M., Pudjianto, D., Djapic, P., Teng, F., Sturt, A., Jackravut, D., Sansom, R., Yufit, V. and Brandon, N. (2012), Strategic assessment of the role and value of energy storage systems in the U.K. low carbon energy future, Technical report, Imperial College London, London, United Kingdom.

Swan, L. G. and Ugursal, V. I. (2009), 'Modeling of end-use energy consumption in the residential sector: A review of modeling techniques', *Renewable and Sustainable Energy Reviews* **13**(8), 1819–1835.

Tabbone, L., Ravalet, E., Durand-Daubin, M., Kaufmann, V. and Cayla, J.-M. (2016), Spatial location of activities and energy consumption of households in France, *in* 'Proceedings of the DEMAND Centre Conference', Lancaster.

Tibshirani, R. (1996), 'Regression shrinkage and selection via the lasso', *Journal of the Royal Statistical Society. Series B (Methodological)* pp. 267–288.

Torriti, J. (2014), 'A review of time use models of residential electricity demand', *Renewable and Sustainable Energy Reviews* **37**, 265–272.

Torriti, J., Hanna, R., Anderson, B., Yeboah, G. and Druckman, A. (2015), 'Peak residential electricity demand and social practices: Deriving flexibility and greenhouse gas intensities from time use and locational data', *Indoor and Built Environment* **24**(7), 891–912.

U.S. Bureau of Labor Statistics (2018), American time use survey: Handbook of methods, Technical report, U.S. Department of Labor, Washington, D.C.

Viegas, J. L., Vieira, S. M., Melício, R., Mendes, V. and Sousa, J. M. (2016), 'Classification of new electricity customers based on surveys and smart metering data', *Energy* **107**, 804–817.

Wickham, H. (2016), *Ggplot2: Elegant Graphics for Data Analysis*, Springer-Verlag, New York.

Wickham, H., François, R., Henry, L. and Müller, K. (2019), 'Dplyr: A grammar of data manipulation'.

Widén, J., Lundh, M., Vassileva, I., Dahlquist, E., Ellegård, K. and Wäckelgård, E. (2009), 'Constructing load profiles for household electricity and hot water from time-use data—Modelling approach and validation', *Energy and Buildings* **41**(7), 753–768.

Widén, J. and Wäckelgård, E. (2010), 'A high-resolution stochastic model of domestic activity patterns and electricity demand', *Applied Energy* **87**(6), 1880–1892.

Wyatt, P. (2013), 'A dwelling-level investigation into the physical and socio-economic drivers of domestic energy consumption in England', *Energy Policy* **60**, 540–549.

Zou, H. and Hastie, T. (2005), 'Regularization and variable selection via the elastic net', *Journal of the Royal Statistical Society. Series B (Statistical Methodology)* **67**(2), 301–320.

5.A Supplemental tables

Predictor	5–9 p.m. (1) β	5–9 p.m.(2) β	9 p.m.–1 a.m. β	1 a.m.–5 a.m. β	5 a.m.–9 a.m. β	9 a.m.–1 p.m. β	1 p.m.–5 p.m. β
Activity categories							
Care for home	0.46***		0.21*		0.22	0.32	0.28
Care for others					0.23	0.49**	
Food	0.16	0.52***	0.25*		0.25		0.31
Other activity							0.37*
Recreation	0.33*		0.25*			0.42*	0.44*
Appliances							
EV				2.23***			
Gas boiler			−0.37*		−0.59**		
Heat pump					0.76		
No. night storage heaters							0.53**
No. power showers			0.29**			0.35*	
No. TVs/computers		0.19***				0.46**	0.24
Tumble dryer					0.28		0.37*
Under floor heating			0.11		0.34		
Dwelling							
Detached home	0.62***	0.56**	0.20		0.17		0.66**
No. rooms in home		0.20**				0.25	0.21
Terraced home	−0.16						
Socio-demographics							
Bill affordability	−0.24*						
Estimated monthly electric bill		0.27***	0.18		0.26	0.25	
No. of occupants	0.30*		0.14		0.20		0.29
Owns cat		0.31	0.36**		0.35*		
Participated in summer	−0.43***						
Participated in spring/autumn			0.34**				
Tariff: Economy 7 or 10				1.73***			
Tariff: Renewable	−0.42**		−0.43**				
Tenure (rent vs. own)					−0.38*		
Constant	6.19	5.66	5.66	4.10	5.28	4.99	4.75
Observations	171	170	172	172	172	170	170
R^2	0.43	0.25	0.43	0.28	0.33	0.29	0.39
Adj R^2	0.40	0.23	0.39	0.27	0.28	0.26	0.35

Note: $*p < 0.05$, $**p < 0.01$, $***p < 0.001$.

Table 5.A.1: Standardised regression coefficients for Table 5.5.

Predictor	5–9 p.m. (1) β	5–9 p.m.(2) β	9 p.m.–1 a.m. β	1 a.m.–5 a.m. β	5 a.m.–9 a.m. β	9 a.m.–1 p.m. β	1 p.m.–5 p.m. β
Activity categories							
Care for home					1.05**	0.44*	0.50**
Care for self					0.21		
Food	0.04	0.43**	0.28**			0.23	0.38*
Recreation	0.40**		0.43***			0.53**	0.44*
Appliances							
EV				2.23***			
Gas boiler			−0.32*		−0.65***		
No. night storage heaters							0.56***
No. power showers			0.34**			0.29	
No. TVs/computers						0.38*	0.20
Tumble dryer						0.34	0.38*
Under floor heating			0.12		0.40*	0.38	
Dwelling							
Detached home	0.69***	0.66**	0.20		0.28	0.18	0.76**
No. rooms in home		0.22				−0.02	0.19
Terraced home	−0.15						
Socio-demographics							
Estimated monthly electric bill		0.34*	0.12		0.29	0.13	
No. of occupants	0.39**		0.19		0.35*	0.28	0.35*
Owns cat			0.33**		0.33		
Participated in summer	−0.43**						
Participated in spring/autumn			0.36**				
Participated on Fri-Sat						0.56*	
Tariff: Economy 7 or 10				1.73***			
Tariff: Renewable	−0.53**		−0.40**				
Constant	6.19	5.72	5.62	4.10	5.34	4.69	4.74
Observations	171	170	172	172	172	170	170
R^2	0.43	0.25	0.43	0.28	0.33	0.29	0.39
Adj R^2	0.40	0.23	0.39	0.27	0.28	0.26	0.35

Note: $^*p < 0.05$, $^{**}p < 0.01$, $^{***}p < 0.001$.

Table 5.A.2: Standardised regression coefficients for Table 5.6.

Chapter 6

Cluster analysis and prediction of residential peak demand profiles using occupant activity data (Paper 3)

Aven Satre-Meloy, Marina Diakonova, Philipp Grünewald
Published in *Applied Energy*

Abstract

Researching the dynamics of residential electricity consumption at finely-resolved timescales is increasingly practical with the growing availability of high-resolution data and analytical methods to characterise them. One methodological approach that is popular for exploring consumption dynamics is load profile clustering. Despite an abundance of available algorithmic techniques, clustering load profiles is challenging because clustering methods do not always capture the temporal aspects of electricity demand and because clusters are difficult to explain without additional descriptive household data. These challenges limit the use of cluster analysis to better understand behavioural and other drivers of electricity demand patterns.

We address these challenges by applying a novel clustering approach to a unique data set of high-resolution electricity and occupant time-use data from U.K. households. We cluster cumulative rather than raw load profiles to capture their full shape. Our clustering approach identifies two distinct patterns of electricity demand during evening weekdays (5–9 p.m.), which are primarily differentiated by the timing of their peak demand. Next, we apply several classification algorithms to assess the potential for using time-use activity data to predict membership in these distinct clusters. The methods we use are suited to this predictive modelling context and are able to identify

key activities driving patterns of electricity demand. We discuss how such an approach can inform more targeted strategies for residential peak demand reduction and response interventions as well as improve our understanding of constraints and opportunities for demand-side flexibility in the residential sector.

6.1 Introduction

The residential sector accounts for around 30% of total U.K. electricity consumption and 50% of U.K. national peak electricity demand (BEIS, 2018; Ofgem, 2010). Reducing overall electricity demand in buildings through energy efficiency is an important component of climate mitigation strategies. But in addition to reducing total demand, better understanding the temporal aspects of electricity demand is important in energy research, especially in the residential sector where temporal variability is high (McLoughlin et al., 2013). Reducing system peak demand lowers the costs and carbon-intensity of electricity generation (National Infrastructure Commission, 2016). Delivering more responsive demand can also help integrate variable renewable energy into existing power systems. Beyond making electricity use more efficient, understanding how to increase the flexibility of use will benefit any low-carbon energy strategy (Grunewald and Diakonova, 2018*b*; Satre-Meloy and Langevin, 2019).

Deeper insight into the factors that influence consumption patterns across hours of the day can inform solutions for demand flexibility and can improve efforts to target consumers for time-sensitive reductions in demand. Recent research suggests occupant activity data may be particularly valuable for understanding patterns of electricity demand during different times of day (Satre-Meloy et al., 2019). Access to high-resolution electricity demand and activity data can strengthen these insights if appropriate methods are used to characterise them.

One technique used to identify temporal variations in electricity consumption that has gained popularity in recent years is cluster analysis of electric load profiles (Albert and Rajagopal, 2013; Granell et al., 2015*a*; Kwac et al., 2018; McLoughlin et al., 2015; Rhodes et al., 2014; Viegas et al., 2016). Cluster analysis is an algorithmic approach used to identify homogeneous groupings of data where no *a priori* grouping exists (Liao, 2005). This data mining technique has gained popularity with the rise of 'big' data and machine learning, and it has been applied in recent research to interval meter data from residential customers.

Cluster analysis of load profiles in the residential sector can be used for numerous practical purposes, including identifying characteristics that correlate with different energy usage patterns (Beckel et al., 2014; McLoughlin et al., 2015), creating more accurate profiles for groups of utility customers (Räsänen et al., 2010; Stephen et al., 2014), which can then be used to design more appropriate tariffs (Flath et al., 2012), or targeting customers that may be particularly viable for DSR schemes (Albert and Rajagopal, 2013; Cao et al., 2013; Kwac et al., 2014). These applications of cluster analysis can aid utility programme designers in discovering potential energy saving opportunities for customers based on their specific patterns of usage while also enabling them to better coordinate demand-side resources for energy management.

Despite the emergence of this research, however, identifying representative classes of customers that can be explained by socio-demographic or other household data remains challenging for several reasons. First, clustering results have been shown to be highly sensitive to numerous methodological considerations and assumptions (Hsu, 2015). Especially in the context of clustering load profiles, comparatively few studies have addressed the issue that cluster analysis of time-series electricity data using conventional distances does not adequately capture temporal variations in consumption patterns (Iglesias and Kastner, 2013; Piao et al., 2014; du Toit et al., 2016).

Second, because clustering is a technique that always yields some segmentation of data into groups, it can be difficult to gauge whether in addition to being mathematically sound, the clustering results are helpful for utilities or policymakers. In the context of targeting customers for DSR schemes, designing better tariffs, or identifying characteristics associated with energy usage behaviours, it is essential that clusters can be explained by associated descriptive data from the dwelling and its occupants.

Numerous studies have highlighted the need for additional information from households to verify the social, lifestyle, and behavioural patterns that underpin temporal variations in electricity consumption (Albert and Rajagopal, 2013; Kwac et al., 2014, 2018; Smith et al., 2012). These studies stop short of investigating the social forces that influence the magnitude and timing of demand because of data availability issues, and they instead make assumptions or simple guesses about the lifestyle and activity factors that explain different clusters (Flath et al., 2012; Haben et al., 2016; Teeraratkul et al., 2018). Empirical, data-based evidence, rather than evidence from models based on assumptions or guesses as to which end uses are driving consumption, is lacking.

This paper addresses these two challenges through an integrated cluster and classification analysis of evening peak-period electricity demand in U.K. households. It introduces a novel data set of high-resolution electricity demand and detailed time-use activity data and applies cluster and classification analyses to these for the purpose of identifying key activities that drive different patterns of electricity use from 5–9 p.m. in U.K. homes. It makes two primary contributions to the literature on household load profile segmentation and prediction.

First, it makes a methodological contribution by introducing a novel approach for pre-processing load profiles prior to clustering that focuses the cluster analysis on capturing temporal variations in electricity demand. We propose clustering cumulative rather than raw load profiles, which enables more appropriate use of Euclidean distances to properly account for temporal differences between clusters. This approach to clustering load profiles does not appear to have been used in previous research. After we apply this pre-processing step, we cluster profiles using several different Euclidean distances and two types of clustering algorithms and then select the best performing methods, as evaluated using several cluster validity indicators (CVIs), to determine the optimal cluster assignment for our data.

The paper's second contribution is its incorporation of detailed time-use activity data from household occupants to predict cluster membership. We apply two classification models, which are specifically chosen for their ability to handle the complexities inherent to incorporating thousands of reported activities in the prediction task, as well as for their ability to identify key activities associated with distinct demand patterns. The aim of this stage of the analysis is to show how access to household time-use data can improve segmentation of customers and identify activities driving their demand. To the authors' knowledge, no previous studies have combined load profile segmentation via cluster analysis with time-use activity data in a predictive context to identify activities driving temporal variations in household electricity demand. This novel approach can inform efforts to develop more targeted and effective strategies for peak demand reduction and response in the residential sector.

The paper is organised as follows: Section 6.2 reviews related work, both on clustering electric load profiles and using time-use activity data in models of building electricity consumption. Section 6.3 describes the data used for subsequent analyses. Section 6.4 describes the clustering and classification analyses undertaken in this paper. Section 6.5 presents results, and Section 6.6 discusses implications of these and avenues for future research.

6.2 Related work

6.2.1 Cluster analysis of electric load profiles

There is a large and active literature on electric load profile segmentation using cluster analysis. In this review, we mostly focus on the empirical findings of papers that have used this approach, but we note that the literature reviewing various methodological approaches and considerations is equally active. The authors recommend Liao (2005), Wang et al. (2015), and Zhou et al. (2013) for reviews of this literature.

We summarise previous empirical findings in Table 6.1. As shown in the table, studies are concentrated primarily in Europe and the U.S., with one study each from China, South Korea, and South Africa. Numerous studies used the same data sets, such as the Irish Commission for Energy Regulation (CER) smart metering trial data. The residential sector is the most commonly analysed. In terms of temporal resolution, most studies use data sampled at either 15 minutes, 30 minutes, or hourly, and data are in most cases sampled for one part of the year only (e.g., summer or fall). Sample sizes for these studies vary considerably. While cluster analysis in other disciplines is usually applied to large data sets, for load profile segmentation there are numerous studies with samples in the hundreds of customers. Several, however, have much larger sample sizes in the hundreds of thousands.

The number of clusters (K) determined as optimal in different studies is highly variable and often dependent on the clustering algorithm used and nature of the data. Still, there does seem to be some consistency across studies with most finding 5–15 clusters. In numerous cases the authors do not select the number given as optimal by a CVI but instead use some judgement to determine a cut-off in number of clusters in light of the segmentation goal or to better match numbers that would be useful for electricity companies (e.g., Figueiredo et al., 2005; McLoughlin et al., 2015). Finally, the right-most columns in Table 6.1 indicate whether and what types of auxiliary household data (HD) are linked to clusters and whether or not these data are used in statistical models (SM) to predict cluster assignment. About a third of studies reviewed take this additional step after clustering.

Study	**Country**	**Building sector(s)**	**Data period and resolution**	**Sample size**	**Summary**[a]		
					K	HD[b]	SM
Albert and Rajagopal (2013)	U.S.	Residential	March-October 2010; 10-minute	952	8	AF, SDF	**x**
Azaza and Wallin (2017)	Ireland	Residential	2009-2010	4232	5		
Cao et al. (2013)	Ireland	Residential	July 2009-December 2010; 30-minute	4000	14		
Chicco et al. (2006)	Romania	Commercial, Industrial	Date not given; 15-minute	234	15		
Dent et al. (2011)	U.K.	Residential	3 months in summer & winter; hourly	93	9		

[a] Optimal number of clusters (K); Links household data (HD); Conducts statistical modelling (SM)
[b] Appliance factors (AF); Dwelling factors (DF); Socio-demographic factors (SDF)

Continued on next page

Study	**Country**	**Building sector(s)**	**Data period and resolution**	**Sample size**	**Summary**[a]		
					K	HD[b]	SM
Figueiredo et al. (2005); Rodrigues et al. (2003)	Portugal	Residential, Commercial, Industrial	3 months in summer & winter; 15-minute	165	6–9		**x**
Flath et al. (2012)	Germany	Residential, Commercial	Winter 2010-2011; 15-minute	215	10–14		
Gouveia and Seixas (2016)	Portugal	Residential	July-August 2014; 15-minute	265	10	AF, DF, SDF	
Granell et al. (2015*b*)	U.K. & Bulgaria	Residential	2010; various temporal resolutions	197	2–4		
Granell et al. (2015*a*)	U.K. & Bulgaria (res), U.K. (com)	Residential, Commercial	April-November 2010 (res), 2009-2010 (com); 1-minute (res), 30-minute (com)	197 (res), 1207 (com)	4 (res), 6 (com)	DF, SDF	**x**
Haben et al. (2016)	Ireland	Residential	2009-2010; 30-minute	3622	10		
Iglesias and Kastner (2013)	Spain	Commercial	August 2011- January 2012; hourly	5	n/a		
Jin et al. (2017)	U.S.	Residential	June 2011-May 2012; hourly	100,000	15–30		
Kwac et al. (2013, 2014, 2018)	U.S.	Residential	April 2008-October 2011; hourly	220,000	n/a		
Liu et al. (2012)	Finland	Residential	2009; quasi-daily	11,964	4	DF, SDF	
McLoughlin et al. (2015)	Ireland	Residential	July-December 2009; 30-minute	3,941	10	AF, DF, SDF	**x**
Panapakidis et al. (2014)	Greece	Commercial	January 2010-December 2011; 15-minute	27	15		
Piao et al. (2014)	South Korea	Residential, Commercial, Industrial	January 2007; hourly	1,205	4		
Räsänen et al. (2010)	Finland	Residential, Commercial, Industrial	2008; hourly	3,989	19		
Rhodes et al. (2014)	U.S.	Residential	November 2012-October 2013; hourly	103	2	AF, DF, SDF	**x**
Smith et al. (2012)	U.S.	Residential	Summer 2011; 15-minute	6,662	6		

[a] Optimal number of clusters (K); Links household data (HD); Conducts statistical modelling (SM)
[b] Appliance factors (AF); Dwelling factors (DF); Socio-demographic factors (SDF)

Continued on next page

Study	**Country**	**Building sector(s)**	**Data period and resolution**	**Sample size**	**Summary**[a]		
					K	HD[b]	SM
Teeraratkul et al. (2018)	U.S.	Residential	July-August 2012; hourly	1,057	12		x
du Toit et al. (2016)	South Africa	Residential, Commercial Industrial	Summer 2007-2014; 30-minute	Not given	2–3		
Tsekouras et al. (2008)	Greece	Industrial	2003; 15-minute	93	8–12		
Viegas et al. (2016)	Ireland	Residential	2009-2010; 30-minute	3,440	4	AF, DF, SDF	x
Xu et al. (2017)	U.S.	Residential	January 2015; 15-minute	Not given	9		
Zhang et al. (2012)	China	Residential, Commercial, Industrial	August 2009; 30-minute	155	5		

[a] Optimal number of clusters (K); Links household data (HD); Conducts statistical modelling (SM)
[b] Appliance factors (AF); Dwelling factors (DF); Socio-demographic factors (SDF)

Table 6.1: Summary of previous electric load profile clustering papers reviewed.

From the papers reviewed, we generally find that cluster analysis for load profile segmentation is done for several different but often related purposes. These include, in order of frequency in the literature reviewed, 1) comparative analyses to determine which combinations of distances and clustering algorithms perform the best for load profiles given varying approaches for validating cluster results or different temporal resolutions for the data (Cao et al., 2013; Chicco et al., 2006; Granell et al., 2015*b*; Iglesias and Kastner, 2013; Jin et al., 2017; Piao et al., 2014; Rodrigues et al., 2003; Teeraratkul et al., 2018; du Toit et al., 2016; Tsekouras et al., 2008; Zhou et al., 2013); 2) studies of stability of cluster assignments or entropy of load shapes for individual households, which measures how often households consume the same daily load shape pattern (Albert and Rajagopal, 2013; Kwac et al., 2013, 2014; Piao et al., 2014; Smith et al., 2012; Teeraratkul et al., 2018; du Toit et al., 2016; Xu et al., 2017; Zhang et al., 2012); 3) analyses that use statistical methods to link load profile clusters to dwelling or other household data (Figueiredo et al., 2005; Gouveia and Seixas, 2016; Granell et al., 2015*a*; McLoughlin et al., 2015; Rhodes et al., 2014; Viegas et al., 2016); 4) investigations of the variability in timing of peak demand, the contributions of different customer segments to peak demand, or related time-of-day and seasonal effects on electricity demand patterns (Azaza and Wallin, 2017; Liu et al., 2012; Smith et al., 2012; Xu et al., 2017).

Given that this paper's aims are primarily related to 3 and 4, several findings specific to these aims from the literature are given here. Studies that have used statistical models to associate household data with cluster assignment show that several dwelling, socio-demographic, and appliance factors are influential in predicting load profile cluster. Granell et al. (2015*b*) find that number of bedrooms and occupants are key drivers, while Gouveia and Seixas (2016) find that dwelling year of construction, floor area, and heating and cooling equipment, along with number of occupants and monthly income, are important. In Viegas et al. (2016), the authors find customer employment status and age are significant, as well as number of dishwashers, electric cookers, and washing machines. McLoughlin et al. (2015) associate dwelling, socio-demographic, and appliance ownership variables to 10 profile clusters, finding that factors such as number of bedrooms, ownership of energy-intensive appliances, and age of head-of-household (HoH) are significant. The authors note, however, that in many cases, standard errors for regression coefficients are large owing to small samples in certain variable sub-categories. Finally, Rhodes et al. (2014) find that variables such as working from home, hours of television watched per week, and education levels have significant relationships with average load profile shape but that these are variable across seasons.

6.2.2 Prediction and classification of building electricity consumption

Cluster analysis is often used as a pre-processing step for predicting or classifying building electricity consumption, and there exist numerous data-driven approaches that can be applied to the clustering results in a predictive context. Given that we combine both cluster analysis and classification in an integrated analysis in this paper, in this section we provide a brief review of the literature on building energy prediction and classification.

Wei et al. (2018) present a recent and detailed review of the main categories of approaches for building energy prediction. These categories are described as white-box, grey-box, and black-box approaches. White-box approaches refer to physics-based models that use detailed building characteristics, such as appliance power ratings, and thermodynamic principles to simulate building operations and predict energy consumption. These approaches do not require historical energy consumption information and have high accuracy but also require detailed data measurements of building characteristics, which makes them more computationally intensive (Swan and Ugursal, 2009). Grey-box approaches require both physical building information as well as historical data and combine statistical methods with building physics models to predict building electricity consumption (Wei et al., 2018). Conditional demand analysis (CDA) is a specific method that performs regression to determine the use level of individual appliances (Parti and Parti, 1980). While these methods benefit from the use of statistical methods and thus can be applied to larger samples, one disadvantage is that they involve complex interactions between end uses, which increases the uncertainty of their predictions. The final category, black-box approaches, refers to methods that predict building electricity consumption strictly using historical data and statistical analysis. These approaches are becoming increasingly popular for their ability to incorporate a wide range of data types, including occupant behaviour and socio-demographic factors, without sacrificing predictive accuracy (Zhao and Magoulès, 2012).

As this paper's classification approach falls under the black box or "data-driven" family of building energy consumption prediction, a short review of several common methods within this family is provided here. Some of the more well-known data-driven methods include various types of regression (e.g., multiple linear regression, polynomial regression) as well as more complex non-linear methods such as support vector machines (SVM), artificial neural networks (ANN), and decision trees (DT). In a systematic review of the data-driven building energy consumption literature, Amasyali and El-Gohary (2018) found that around half of the studies reviewed employ ANNs, a quarter use either SVM or regression methods, and only around 4% use decision trees. Due to their highly non-linear nature, both ANNs and SVMs and their extensions are able to achieve high accuracy in energy consumption prediction tasks (e.g., Edwards et al., 2012). Regression approaches have historically showed lower predictive accuracy than ANNs and SVMs because of their inflexibility (Wei et al., 2018), but recent extensions to standard multiple linear regression, such as regularization methods, deliver performance gains in terms of prediction and model interpretability (Satre-Meloy, 2019; Hsu, 2015). DTs have been shown to perform as well as ANNs and SVMs in some contexts (Tso and Yau, 2007; Yu et al., 2010), though in general their predictive accuracy is somewhat lower (Wei et al., 2018).

There are many other data-driven algorithms for building energy consumption, and a more detailed review is beyond the scope of this paper. We recommend both Wei et al. (2018) and Amasyali and El-Gohary (2018) for more in-depth reviews of these approaches. As will be described in Section 6.4, we apply two of the four types of data-driven methods reviewed above—regression and decision trees. Our approach, however, goes beyond previous research by applying extensions to these two methods, which are necessary given the complexities inherent in our predictive modeling context.

6.2.3 Time-use activity data in electricity demand models

This paper uses time-use data to predict load profile cluster membership, so to familiarise readers with this area of research, in this section we provide a brief review of time-use data and its application to energy demand modelling. For a more complete review, we direct interested readers to Torriti (2014).

Time-use activity data have been collected as part of national studies on how people spend their time (Gershuny, 2008; U.S. Bureau of Labor Statistics, 2018). More recently, these data have been used as inputs to energy demand models (McKenna and Thomson, 2016; Richardson et al., 2010; Torriti et al., 2015; Widén et al., 2009; Widén and Wäckelgård, 2010). Studies investigating the associations between time-use and electricity use have found some consistency in terms of the activities that influence electricity use patterns. Key activities linked to higher demand include meal preparation-related activities and meal times (Anderson and Torriti, 2018; De Lauretis et al., 2017; Druckman et al., 2012; Jalas and Juntunen, 2015), housework and personal time (Palmer et al., 2013; Widén et al., 2009), and some recreational activities, especially TV use (Satre-Meloy et al., 2019). Activities that are linked with lower demand unsurprisingly include sleeping and resting (Druckman et al., 2012; Jalas and Juntunen, 2015).

Approaches for incorporating time-use data into energy demand models vary widely, and a detailed review is beyond the scope of this paper. Key limitations, however, have been summarised in previous literature (McKenna et al., 2017; Torriti, 2014). Among these, several are particularly relevant for this paper. First, as McKenna et al. (2017) note, time-use studies were not originally designed for energy demand modelling, so time-use data often are not tailored to questions regarding electricity use. For instance, study participants are not asked to differentiate between energy-intensive or low-energy alternatives of the same activity (e.g., cold meals versus hot meals). Second, as time-use data have historically been collected using paper diaries that ask participants to write a primary and secondary activity during each 10-minute window of the day (the 'time budget' approach), overlapping activities were not easily accounted for, and some electricity-relevant activities that may occur during shorter time periods but might still be important for understanding demand (e.g., boiling the kettle) were underreported. Furthermore, as Torriti (2014) notes, another key challenge is using activity data to model multiple-person households (cf. Grunewald and Diakonova, 2018*a*).

The data collection process described in the following section was designed with these limitations in mind and addresses these through a new approach to collecting time-use data.

6.3 Data

The collection process for the data analysed in this paper has been described in previous work (Grunewald and Diakonova, 2018*a*; Satre-Meloy et al., 2019). We briefly review the study design and data collection methods and then describe several updates to the data and specific sampling considerations for this paper's analysis.

Data are collected as part of a five-year study for which data collection is ongoing (Grünewald and Layberry, 2015). Participating households, which are recruited online, via e-mail and social media, complete a household survey before participating wherein they provide socio-demographic data along with physical dwelling characteristics and household appliance ownership details. Fuel type used for heating and cooking appliances is also collected here. Households then receive a parcel prior to their selected study date containing an electricity recorder, activity recorder(s), and an instruction booklet. Any household member above the age of eight can participate.

On their selected date, which can be any day of the week excluding holidays, participants are instructed to attach the electricity recorder to their mains electricity. The recorder collects electricity demand data at one second resolution for 28 hours, from 5 p.m. until 9 p.m. on the following day, thus covering two typical peak demand periods in the early evening. The U.K. is winter-peaking, and around two-thirds of data are therefore collected between September and April.

Activities are recorded using a purpose-built app that comes pre-installed on individual devices. The app guides participants through screens where they can enter the activity location, details of the activity, number of people participating, and enjoyment of the activity. A schema representing an example activity entry sequence is shown in Figure 6.1. The 'Activity' screens are tailored to collect numerous details on the type of activity, appliances used, and other energy-relevant details. These screens enable participants to record more detail than would be possible using paper-based activity diaries, which have been used extensively for time-use research (eurostat, 2009). Participants are encouraged to record activities in real time, and activities are recorded at points in time rather than for durations in time. The app also provides for recording activities retrospectively and in the future. These capabilities facilitate reporting of overlapping activities as well as activities that may not fit into the typical 10-minute reporting windows used in past time-use studies. Previous analyses have examined the accuracy of activity reporting in this data set and have found around 80% of reports for certain activities are entered within 10 minutes of the activity itself (see Grunewald and Diakonova, 2018*a*). A description of the data storage and handling procedures is given in Grünewald and Diakonova (2019*a*).

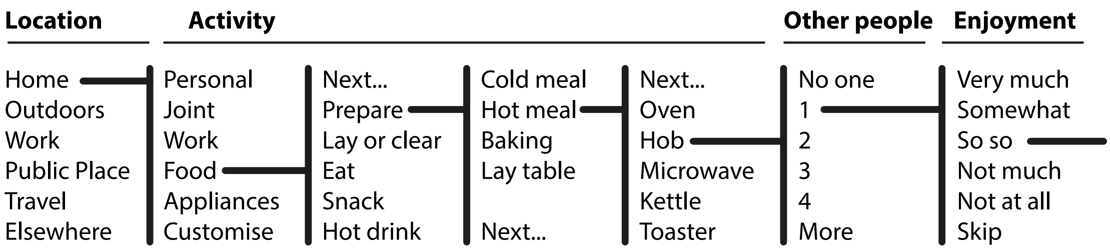

Figure 6.1: Example of an activity entrance sequence on the activity recorder (Grünewald et al., 2017).

In addition to the procedures mentioned above, we exclude households with the following from the sample:

- poor electricity readings (e.g., failure to attach the recorder);
- missing electricity readings between 5 and 9 p.m. Multiple imputation is unnecessary due to the low incidence of missing data ($N = 14$, $\sim$5% of sample);
- no activities reported between 5 and 9 p.m.;
- PV (unreliable net-demand electricity readings);
- weekend participation dates (weekend demand tends to be lower and different to weekday demand (e.g., Figueiredo et al., 2005)).

6.4 Methods

In this section, we describe our methods for clustering 5–9 p.m. load profiles and then predicting load profile cluster using occupant activity data. The overall analysis approach taken in this paper is most similar to the approaches demonstrated in McLoughlin et al. (2015); Rhodes et al. (2014); Viegas et al. (2016), and Figueiredo et al. (2005). The key difference is our use of occupant activity data for the predictive modelling task and our application of regularisation methods for identifying key predictive activities. A graphical overview of our methodological approach is shown in Figure 6.2 and explained step-by-step in Sections 6.4.1–6.4.3. Section 6.4.1 describes how we approach the clustering task and then pre-process and normalise the data, including the assumptions we make at these steps. Section 6.4.2 explains our procedure for validating cluster results and assigning cluster membership. Section 6.4.3 then describes the predictive modelling component of our analysis, which uses occupant activity data to predict cluster membership. In this section we also describe our model training and evaluation procedures.

6.4.1 Data pre-processing and normalisation

In this paper we take a raw-data-based approach (McLoughlin et al., 2015; Räsänen et al., 2010; Rhodes et al., 2014; Teeraratkul et al., 2018; Xu et al., 2017) and apply clustering algorithms to whole time-series electricity data rather than extracted features (e.g., Azaza and Wallin, 2017; Haben et al., 2016; Kwac et al., 2018). We downsample our raw electricity readings to create hourly profiles[1] where the hourly average is assigned to the interval beginning hour (e.g., electricity use data from 5:00 to 5:59 p.m. are averaged and assigned to 5:00 p.m.). We then extract the four values corresponding to the hourly averages from 5 to 8 p.m. for day 1 and day 2 of the study period so that each household can be represented by two 4-element vectors, given by:

[1] We use hourly averages because we do not find clear performance gains in the clustering task when sampling our data at higher resolutions, such as 30-minute or 15-minute (cf. Granell et al., 2015*b*), and because using hourly averages reduces the dimensionality of the data considerably.

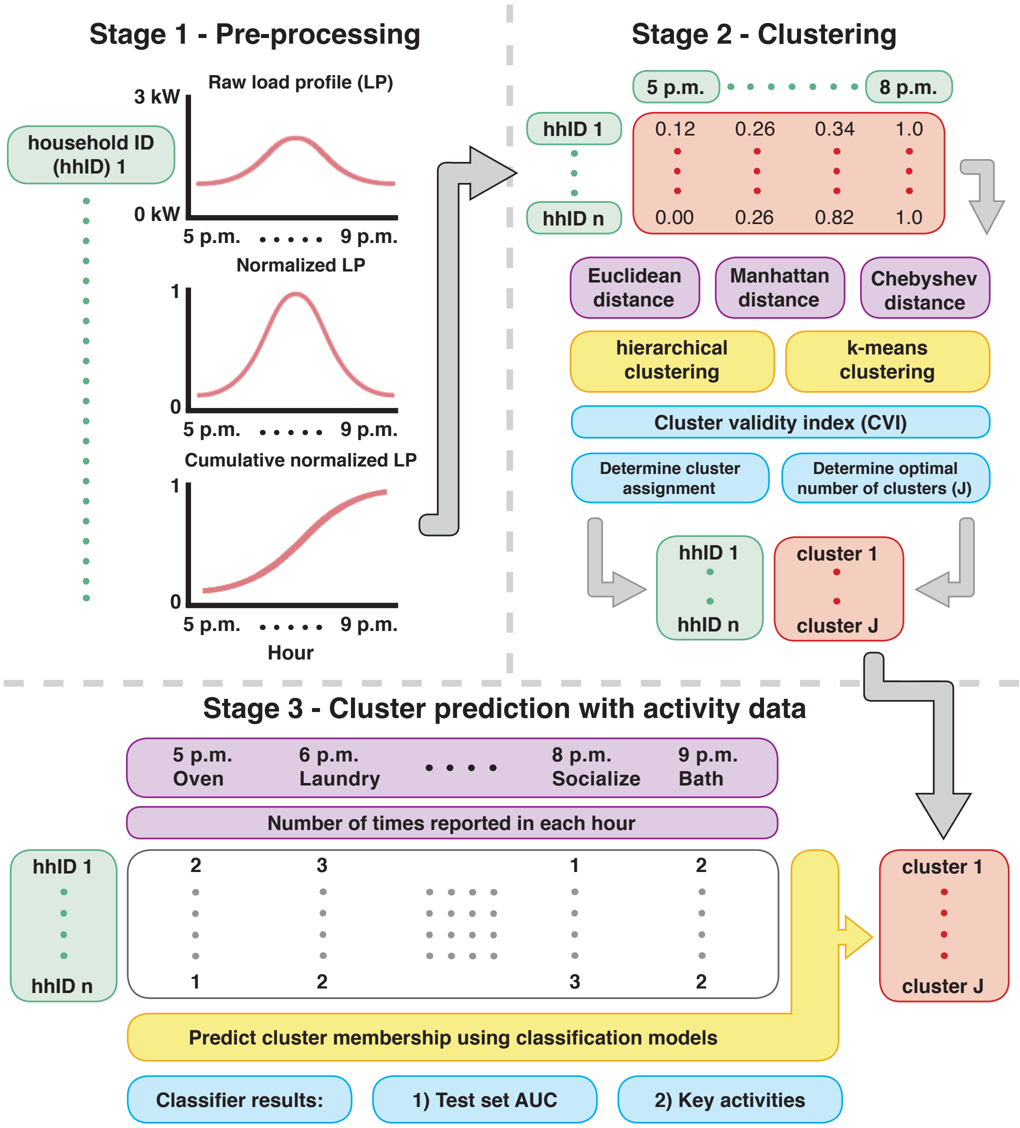

Figure 6.2: Overview of methodological approach for clustering and classifying load profiles, with Stages 1, 2, and 3 described in Sections 6.4.1–6.4.3, respectively.

$$l_i = \begin{pmatrix} P_{d5} \\ P_{d6} \\ P_{d7} \\ P_{d8} \end{pmatrix} \tag{6.1}$$

where l_i is the load profile for each home i, and P_{dh} is the hourly power demand in Watts (W) of home i on day d at hour h.

The two load vectors per household (day 1 and day 2) are treated as independent in this analysis in order to increase the sample size. Figure 6.A.1 in the appendix shows that when only day 1 profiles are analysed, results for cluster membership and shape are similar to those obtained on the combined sample. This step improves the performance of both the cluster and classification algorithms we use, as these are typically better suited to larger sample sizes.

Our final sample size consists of $N = 269$ household load profiles. The normalisation technique we use is that proposed by Jin et al. (2017) to normalise profiles by first subtracting minimum demand from each hourly value ('de-minning') and then dividing each hour's demand by the 'de-minned' total. For our 5–9 p.m. profiles, we subtract the minimum over this four-hour period rather than the daily minimum. After normalisation, each household load profile represents hourly discretionary demand. We use this normalisation technique because our focus is segmenting households by the temporal variation in their electricity use in order to identify different shapes of peak-period demand profiles. We expect de-minning is especially appropriate for linking occupant activities to electricity demand patterns, as activities are most likely drivers of discretionary changes in usage (as opposed to baseload demand). De-minning also mitigates distortion of load profiles for households with high baseload (Jin et al., 2016).

We take one final data-preprocessing step to more appropriately capture the temporal variations in de-minned hourly load profiles during clustering. When clustering time-series data, it is important to account for their ordered, temporal nature. Even though distances in Euclidean space are used widely for clustering time-series data, these do not capture temporal variations and thus are inadequate for determining dissimilarity of time-series vectors. A simple example of this issue is given in Figure 6.3, which plots three 6-element vectors with magnitude 5 as raw time series. Using a Euclidean distance to calculate the pairwise distance for each pair of these vectors results in the same distance equal to $500\sqrt{2}$. In other words, when using a Euclidean distance, these vectors are equidistant from one another. However, given their overall shape, we would likely say profiles a and b are more similar to one another than they are to profile c.

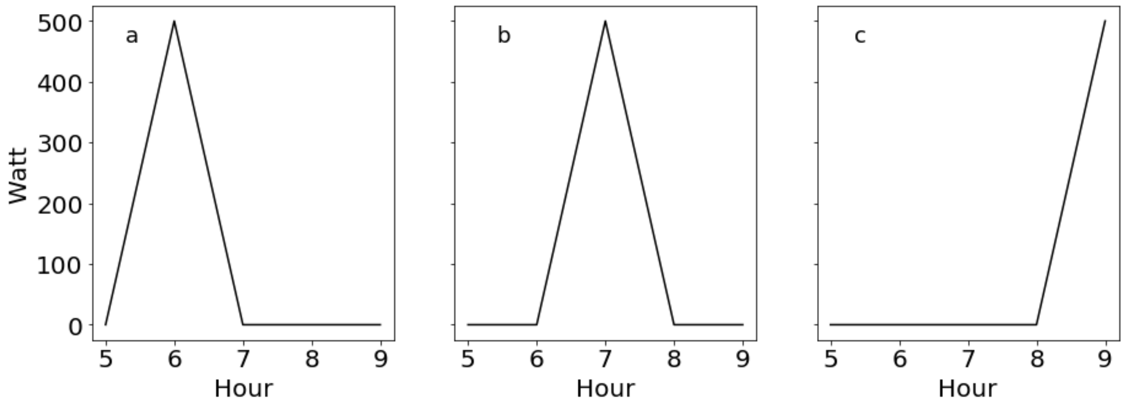

Figure 6.3: Example of three hypothetical load profiles with equal Euclidean distances between each pair.

Several studies have noted the issue of Euclidean distances not appropriately handling time-series load data (Teeraratkul et al., 2018; du Toit et al., 2016). We propose a novel solution to this problem: instead of clustering the raw time-series data, we take the integral of the load shape l_i, constructing a new sequence C_i of the same length that gives the cumulative load, where each component is of the form:

$$C_i = \sum_{h=5}^{8} P_h \tag{6.2}$$

where C_i is the cumulative load profile for each home i, and P_h is the hourly electricity demand (W) of home i at hour h. Returning to the example above, the new cumulative load profiles shown in Figure 6.4 are no longer equidistant. The Euclidean distance between d and e is 500, the distance between d and f is 1000, and the distance between e and f is $500\sqrt{3}$. Clustering the cumulative load profiles would thus group together d and e, which better accounts for the similarities in their temporal shape.

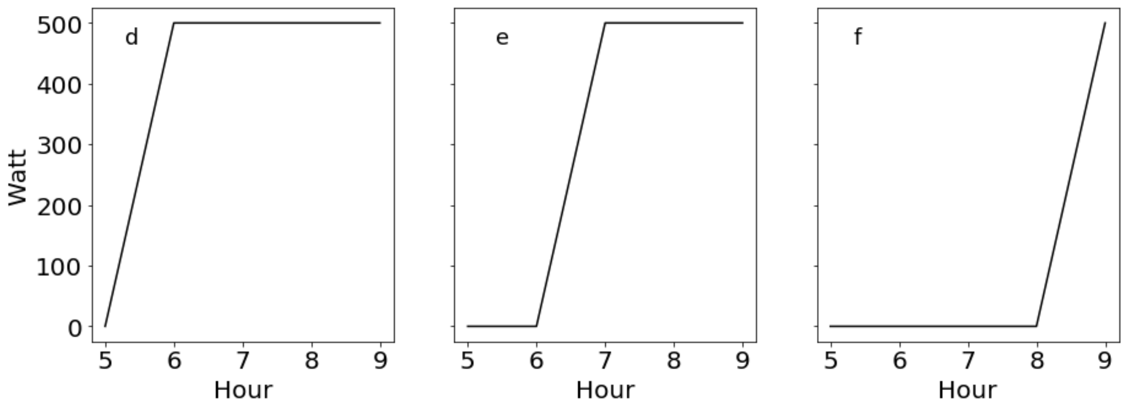

Figure 6.4: Cumulative load profiles with Euclidean distances such that d and e are closer together than either is to f.

While several other approaches have been proposed for clustering time-series data, such as the specific distance measures proposed in Liao (2005), we argue our approach is advantageous both in its conceptual simplicity as well as its ease of implementation. Based on our review of existing load profile cluster analyses, this step is wholly absent from past work. Xu et al. (2017) suggest a similar approach to integrate load shapes, but they do so to consider maximum demand over the time series rather than to capture the temporal shape of load profiles.

After these pre-processing and normalisation steps, we next move to clustering these de-minned, cumulative load profiles where we are able to achieve the simultaneous goals of focusing on temporal variations in discretionary demand while also enabling straightforward use of Euclidean distances.

6.4.2 Cluster analysis of peak-period load profiles

In this section we describe our approach to cluster peak-period, de-minned, cumulative load profiles. Because numerous studies have identified the issue of unstable clustering results when using just one clustering algorithm to determine the optimal number of clusters (Hsu, 2015; Zhou et al., 2013), we address this issue by using the *NbClust* package in R to compare clustering results across numerous distance and algorithm combinations. Specifically, we compare three measures of Euclidean distance, two types of clustering algorithm, and six cluster validity indices (CVIs). Our aim is to ensure the cluster assignment is robust to numerous methodological considerations (cf. Iglesias and Kastner, 2013; Tsekouras et al., 2008).

The distance measures used are three variants of the Minkowski metric, which is defined as:

$$D(X,Y) = \sqrt[p]{\sum_{i=1}^{n} |(x_i - y_i)^p|} \tag{6.3}$$

where $D(X,Y)$ gives the distance of order p between two points $X = (x_1, x_2, \ldots, x_n)$ and $Y = (y_1, y_2, \ldots, y_n) \in \mathbb{R}^n$. The simple Euclidean distance is given when $p = 2$, the Manhattan or city block distance when $p = 1$, and the Chebyshev or maximum distance when $p = \infty$. In addition to our motivation to use Euclidean distances given the pre-processing step taken to cluster cumulative load profiles, these distance measures are some of the most frequently used in previous research and have shown improved clustering results in terms of stability and performance across various CVIs (Chicco et al., 2006; Iglesias and Kastner, 2013; Jin et al., 2017; du Toit et al., 2016).

We include two types of clustering algorithms in our comparative analysis. These are k-means and several variants of hierarchical agglomerative clustering (hac). From the studies reviewed, these were the most commonly applied, and numerous studies also concluded that these perform well for load profile clustering (Chicco et al., 2006; Granell et al., 2015*a*; Jin et al., 2017; Rodrigues et al., 2003; Tsekouras et al., 2008). For the hierarchical clustering, we use several linkage criteria, which determine how clusters are merged together at successive steps. A list of the algorithms used, along with their formulas

Clustering algorithm	Equation	Description
k-means	$\arg\min_{S} \sum_{k=1}^{J} \sum_{x \in S_k} d\left(x, \mu_k\right)^2$	This method minimises the distance between each household load profile x and the groups μ_k.
Single (hac)	$\min_{ij} d(X_i, Y_j)$	This method minimises the distance between the closest members of the two clusters.
Complete (hac)	$\max_{ij} d(X_i, Y_j)$	This method minimises the distance between the most dissimilar or furthest apart members of the two clusters.
Average (hac)	$\frac{1}{kl} \sum_{i=1}^{k} \sum_{j=1}^{l} d(X_i, Y_j)$	This method minimises the average distance between all pairs of members of the two clusters.
Centroid (hac)	$\lvert\lvert c_A - c_B \rvert\rvert$	This method finds the mean vector location for each of the clusters and takes the distance between these two centroids.
Ward (hac)	$\Delta(A, B) = \sum_{i \in A \cup B} \lVert \vec{x}_i - \vec{m}_{A \cup B} \rVert^2$ $- \sum_{i \in A} \lVert \vec{x}_i - \vec{m}_A \rVert^2 - \sum_{i \in B} \lVert \vec{x}_i - \vec{m}_B \rVert^2$ $= \frac{n_A n_B}{n_A + n_B} \lVert \vec{m}_A - \vec{m}_B \rVert^2$	This method minimises the total within-cluster sum-of-squares. Clusters are combined such that their merger results in minimum information loss.

where:

- for k-means, the centroid of cluster S_k is its mean vector $\mu_k = \frac{1}{|S_k|} \sum_{x \in S_k} x$,
- for single, complete, and average, $(X_1, \ldots, X_k)$ and $(Y_1, \ldots, Y_j)$ are the observations for clusters X and Y,
- for centroid, c_A and c_B are the centroids of clusters A and B, respectively,
- for Ward, $\vec{m}_j$ is the center of cluster j, and n_j the number of points in it. Δ is the merging cost of combining A & B.

Table 6.2: Algorithms used to cluster peak-period load profiles.

and descriptions, is given in Table 6.2.

For each combination of distance measure and clustering algorithm, we evaluate the clustering results using several CVIs to determine the optimal number of clusters and cluster assignment. There are many different CVIs reported in the clustering literature, and for load profile clustering specifically, there are several that are used quite frequently. A list of the CVIs we validate our clustering results with is given in the appendix (Table 6.B.1).

To determine the optimal number of clusters, we again use the *NbClust* package, which returns the optimal number as selected by each of the combinations of distance and algorithm. We constrain the number of clusters to between 2 and 20 and do not evaluate results for larger numbers of clusters given our sample size and segmentation objective. While a comparative analysis of clustering results across the different indices (in terms of cluster assignment and optimal number of clusters) is interesting from a methodological perspective, doing so is beyond the scope of this paper.

Instead, to determine the optimal number of clusters and then assign cluster membership, we select the CVI that shows the least variability in index scores across the combinations of distances and algorithms and for varying numbers of clusters. We take this approach because some indices show widely varying results in index scores across different methods. A less variable index suggests clustering results are robust to different algorithms and distances. To assign final cluster membership, we use the most common clustering assignment across the different distances and algorithms, given the selected CVI and the number of clusters shown to be optimal for the data.

6.4.3 Classification of cluster membership using occupant activity data

The final stage of our analysis aims to train models to predict the cluster membership of a new household, as determined by its peak-period electric load profile, given the activities reported by its occupants. For this analysis, we train two different classifiers. The first is a linear binary logistic regression classifier with elastic net regularisation, and the second is a non-linear random forests classifier. Each of these is described in Section 6.4.3. Before we introduce these methods, however, we first explain how a generic classification algorithm works and how we prepare our activity data for model training.

Classification model set-up

A classifier is a function f that maps a set of predictors p of an observation, in this case a set of activities associated with a given load profile, to a categorical response variable y, which represents one of the J load profile clusters. It is given by Eq. 6.4.

$$f : \Re^p \mapsto y \tag{6.4}$$

$$y \in \{c_1, c_2, \ldots, c_j\}$$

Previous studies that have trained classifiers to predict load profile cluster membership have used various predictors or features, such as household socio-demographic data, physical characteristics of the dwelling, appliance ownership, other occupant-related survey data, and extracted features from load profiles (Albert and Rajagopal, 2013; Granell et al., 2015*a*; McLoughlin et al., 2015; Rhodes et al., 2014; Teeraratkul et al., 2018; Viegas et al., 2016). To the author's knowledge, this is the first paper that attempts to predict load profile cluster membership using occupant time-use data.

We set up our training data in the following way: first, we group individual time-use codes (TUCs) into categories of activities that are similar. For instance, while there are separate TUCs for 'Screen time' to indicate whether a TV, tablet, computer, or other device is used, we group these into one category. Doing so has the benefit of reducing 113 unique TUCs into 61 activity groupings. In total, $N = 3038$ activities are reported in our sample of households during the hours of 5–9 p.m. (across both the first and second days of each household's study participation).

Next, for each household and each hour from 5–9 p.m., we count the number of times an activity is reported. In doing so, we create an 'hour-activity' predictor variable for each household to indicate the number of times a given activity group is reported each hour. We aggregate activities at the hour level to prevent the predictor matrix from getting too large. Even still, given the number of hours and activity groupings, the matrix $\Re^p$ has dimensions 269×177, where the number of weekday load profiles is $N = 269$ and the number of 'hour-activity' predictors is $N = 177$. Figure 6.5 gives activity counts for all activities by the hour in which they are reported.

This approach to the predictive modelling task presents some specific analytical challenges. First, the predictor matrix is relatively wide, and although it is not 'high-dimensional' in the sense of $p \gg n$, the inclusion of such a large number of predictors in a conventional linear model creates challenges for interpretability as well as makes the model more prone to overfitting (Friedman et al., 2001). Second, because the predictors are activity counts and

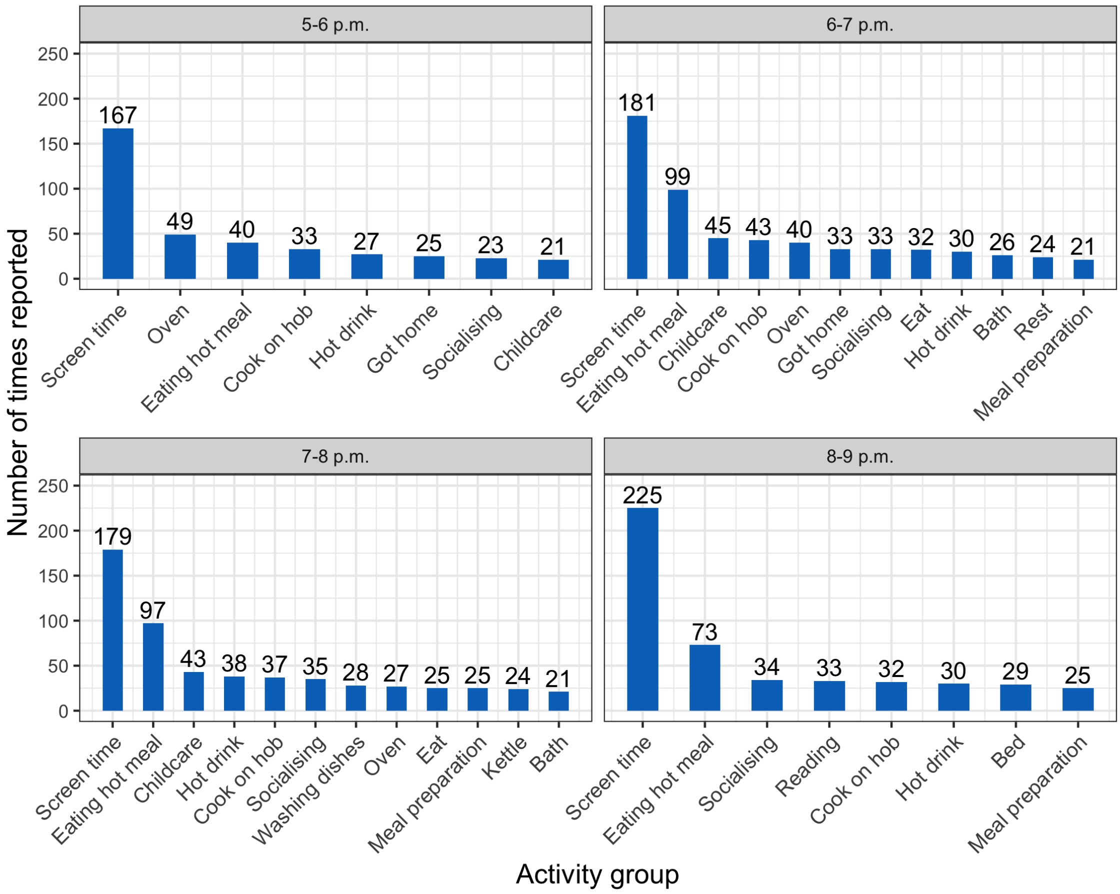

Figure 6.5: Frequency histogram of reported activity groups by hour of evening peak period. Activities reported fewer than 20 times are excluded for brevity.

because many activities are not commonly reported, the predictor matrix is also very sparse with many zero values. This issue raises to two additional challenges, the first of which is that in a two-class modelling scenario with wide data, the classes will almost always be linearly separable, meaning the outcome variable separates a predictor or combination of predictors completely or quasi-completely (Albert and Anderson, 1984). When this is the case, typical linear modelling approaches, such as logistic regression, cannot be used. The second related challenge is that of multicollinearity, which occurs when two or more predictors are highly correlated (Farrar and Glauber, 1967). Multicollinearity can cause issues of stability, where simple changes to the model or data lead to large changes in model parameters. As was the case with the former issue, the predictor matrix as it is designed in this analysis is prone to issues of multicollinearity given some activities are rarely reported and thus several hour-activity variables have near zero variance.

The specific classification methods described in the following sections are designed to address these challenges. We choose a linear method and a non-linear method to compare their performance in correctly classifying load profiles as belonging to clusters given activities reported by household occupants. Our objective in this part of the analysis is both examining the predictive power of activities while also identifying key activities that explain varying patterns of peak-period electricity demand.

Logistic regression with elastic net regularisation

Logistic regression is a Generalised Linear Model (GLM), a generalisation of linear regression that allows for responses that are not continuous, quantitative variables. In the classification setting, especially when the dependent variable is binary, logistic regression is commonly used, and it has been applied to load profile classification in several previous studies (McLoughlin et al., 2015; Rhodes et al., 2014; Viegas et al., 2016). We use logistic regression to determine the likelihood a load profile will be classified as belonging to a specific cluster given the activities occupants report each hour from 5–9 p.m.

In the case of a binary response coded as $Y \in \{0, 1\}$, linear logistic regression models the log-likelihood ratio as a linear combination of predictor variables, as in

$$\log \frac{\Pr(Y = 1|X = x)}{\Pr(Y = 0|X = x)} = \beta_0 + \beta^T x \tag{6.5}$$

where $X = X_1, X_2, \ldots, X_{\mathrm{p}}$ is a vector of predictors, $\beta_0 \in \mathbb{R}$ is an intercept term, and $\beta \in \mathbb{R}^p$ is a vector of parameters or regression coefficients. The inverse of this transformation gives the conditional probability

$$\Pr(Y = 1|X = x) = \frac{e^{\beta_0 + \beta^T x}}{1 + e^{\beta_0 + \beta^T x}} \tag{6.6}$$

which, without restricting the parameters (β_0, β), necessarily constrains the probabilities to lie in [0, 1]. The logit transformation (Eq. 6.5) transforms this conditional probability to a log-linear scale, which enables fitting the model parameters without constraints.

Fitting a logistic regression model is typically done by maximising the likelihood, which is equivalent to minimising the negative log-likelihood

$$\underset{\beta_0,\beta}{\text{minimize}} \left\{ -\frac{1}{N}\mathcal{L}\left(\beta_0, \beta; \mathbf{y}, \mathbf{X}\right) \right\} \tag{6.7}$$

where $\mathbf{y}$ is the N-vector of outcomes (in a binary classification problem, $\mathbf{y} \in \{0, 1\}$), $\mathbf{X}$ is the $N \times p$ predictor matrix, and $\mathcal{L}$ is the log-likelihood function, as given in Eq. 6.5. The negative log-likelihood for logistic regression thus takes the form

$$-\frac{1}{N}\sum_{i=1}^{N}\{y_i \log \Pr\left(Y = 1|x_i\right) + (1 - y_i) \log \Pr\left(Y = 0|x_i\right)\}, \tag{6.8}$$

$$= -\frac{1}{N}\sum_{i=1}^{N}\left\{y_i\left(\beta_0 + \beta^T x_i\right) - \log\left(1 + e^{\beta_0 + \beta^T x_i}\right)\right\}.$$

The model parameters or coefficients (β) in binary logistic regression are interpreted as the expected change in the $\log(odds)$ of $Y = 1$ for each one-unit increase in the corresponding predictor variable, holding other predictor variables in the model constant. To ease interpretation of coefficients, we convert the $\log(odds)$ to odds by finding the exponent Exp$[\beta]$, and we can then find the increase or decrease in the likelihood of being in a particular load profile cluster given an increase in the number of activities reported in each hour of the evening peak period. Model coefficients for the logistic regression are reported both as $\log(odds)$ (β) and odds (Exp$[\beta]$) for this reason. While we have not discussed specific details here, logistic regression can be extended to multi-class classification problems, as well.

Returning to the challenges inherent to our analysis approach, we address these through the application of regularisation to the logistic regression model. Regularisation, also called penalised regression, refers to a family of methods that regularise or constrain the parameter estimation process by adding a penalty term that shrinks the magnitude of model coefficients toward zero. There are several benefits to penalising the estimation process in this way.

The first benefit is that doing so can significantly reduce the variance of regression coefficient estimates. Variance refers to the amount by which model coefficients change if they are estimated with different observational data. Ideally, these estimates should not change much between different data sets, but if a method has high variance, then small variations in the training data can result in large changes in model estimates. This is called overfitting, which often occurs when a very complex model is fit too specifically to the data it is trained on and thus does not generalise well to new data. Especially when p is large, fitting the full model typically increases both its complexity and the variance of its predictions.

In addition to the benefit of reducing the variance of estimates, a certain class of regularisation methods have the benefit of sparsity, meaning they can estimate some of the coefficients to be exactly zero. In this way, these sparse statistical modelling methods can also perform variable selection. Sparsity in statistical models improves their interpretability and is especially useful in the case of large predictor sets. For more detail on statistical learning with sparsity, we direct interested readers to Hastie et al. (2015), and for an overview of regularisation methods and the benefits of applying these in statistical models of energy consumption, see Satre-Meloy (2019).

While there are numerous types of sparse regularisation methods, in this analysis we use the elastic net penalty, introduced by Zou and Hastie (2005). The elastic net penalty is a combination of two other regularisation methods, the lasso and ridge penalties. Its lasso-related properties are the ability to set coefficients exactly to zero, thus yielding sparse models, but it also has properties that make it suitable both for handling data where classes are linearly separable and where there exist highly correlated variables (Hastie et al., 2015). The elastic net penalty is given by

$$\lambda \sum_{j=1}^{p} \left\{(1-\alpha)\beta_j^2 + \alpha\,|\beta_j|\right\} \tag{6.9}$$

where λ is a tuning parameter that determines the magnitude of the penalty that is applied and α is tuned to combine lasso and ridge penalties. If $\alpha = 0$, the coefficients are penalised by the sum of their squares (ridge or ℓ_2 penalty), and if $\alpha = 1$ they are penalised by the sum of their absolute value (lasso or ℓ_1 penalty). Setting α between these gives the elastic net penalty. When applied in the case of logistic regression, the general form of the penalised optimisation problem for minimising the negative log-likelihood is

$$\underset{\beta_0,\beta}{\text{minimize}} \left\{ -\frac{1}{N}\mathcal{L}\left(\beta_0, \beta; \mathbf{y}, \mathbf{X}\right) + \lambda \sum_{j=1}^{p} \left\{(1-\alpha)\beta_j^2 + \alpha\,|\beta_j|\right\} \right\} \tag{6.10}$$

where the notation is the same as in Eqs. 6.7 and 6.9 and the likelihood function $\mathcal{L}$ is as in Eq. 6.8. The elastic-net problem is convex and can be solved with numerous different algorithms. The R package *glmnet* uses coordinate descent and specifically a proximal-Newton iterative approach, which uses a quadratic function to repeatedly approximate the negative log-likelihood (Friedman et al., 2010). The predictors and response are centered and standardised before the algorithm is run.

Because both α and λ can be tuned across a range of values to find the best combination and magnitude of ℓ_1 and ℓ_2 penalties that minimises the training error, cross-validation is typically used to find these. More details on this procedure are given in Section 6.4.3.

One final note on the use of regularisation in logistic regression concerns the question of inference. We do not present typical inferential concepts such as confidence intervals and significance levels for model coefficients because these are not generally appropriate for use with regularisation methods (Taylor and Tibshirani, 2015). There is, however, a growing literature on post-selection inference (Hastie et al., 2015; Lee et al., 2016; Lockhart et al., 2014), which aims to develop methods that can construct confidence intervals and significance tests for regularisation methods. We do not take the additional step of significance testing in this paper given the non-random nature of our study sample.

Random forests

Tree-based methods for classification involve splitting the predictor space into regions that can be used to make predictions for new observations. These methods are so-called because they incorporate a set of rules used to partition the predictor space that can be summarised in a tree. Tree-based methods have been used in many instances for classification and regression related to load profile prediction (Figueiredo et al., 2005; Ma and Cheng, 2016; Viegas et al., 2016; Wang et al., 2018).

Decision trees are a popular statistical learning method because of their simplicity and intuitiveness, but they are often prone to overfitting, which makes them less accurate for predictions. An ensemble method called random forests, which involves producing many decision trees from bootstrap aggregated samples of the data (called bagging) and then giving a consensus prediction across these, can greatly reduce the variance of predictions. Random forests is also able to handle high-dimensional data and is robust to outliers and imbalanced classes (Biau, 2012). The improved predictive performance using random forests over a single decision tree, however, comes at the expense of some interpretability, as it can be difficult to determine which predictors are most important in the model.

In the classification setting, a decision tree predicts each observation to belong to the most commonly occurring class of observations in a given region or node. Because most predictors cannot perfectly separate or predict the correct class for a given observation, they are considered 'impure'. To determine which separation is best given different predictors, a measure of impurity is used. While there are numerous ways to compare impurities, a popular measure is the Gini index (Raileanu and Stoffel, 2004), which is a measure of variance across K classes and is given by

$$G = \sum_{k=1}^{K} \hat{p}_{mk} \left(1 - \hat{p}_{mk}\right) \tag{6.11}$$

$$\hat{p}_{mk} = \frac{1}{N_m} \sum_{x_i \in R_m} I\left(y_i = k\right)$$

where m is a node representing region R_m with N_m observations, and observations in node m are classified to class $k(m) = \arg\max_k \hat{p}_{mk}$, the majority class in node m. The Gini index takes on a small value when a node contains mostly observations from one class. At each split, the predictor with the lowest impurity is used to partition the predictor space until impurity does not decrease further with any subsequent split.

Because random forests works by resampling B bootstrapped training sets with replacement from the training data and then constructing B decision trees for each of these training sets, it is possible that the decision trees could be highly correlated, especially if there is a particularly strong predictor that is selected as the first (or root) node. For this reason, at each split in the tree a random subset of predictors rather than the full set is considered. The size of the random subset M of predictors p that can be selected at each split in the tree is typically chosen such that $M \approx \sqrt{p}$, but this is also a hyperparameter that can be tuned using cross-validation.

Finally, as mentioned above, interpretability is sometimes challenging when using random forests given the bagging procedure. However, we can still obtain an overall summary of the importance of each predictor using the Gini index by summing the total amount the Gini index decreases when split by a given predictor, averaged across the total number of B trees. The variable importance value for each predictor can then be normalised to a range of $[0, 1]$ for comparison.

Model training and evaluation procedure

Models are trained on a sample of 70% of the observations, which we subsequently refer to as the 'training set'. We hold out 30% of our data as the 'test set'. These splits are chosen to create a slightly larger test set than would be the case if we used a standard 80/20 split. As both classifiers require tuning of hyperparameters, we use k-fold cross-validation to tune and train the models. K-fold cross-validation involves splitting the training data into k subsets of equal size. Then each of the k subsets is left out and the other $k-1$ subsets are used to train the model. The left out subset is what the trained model predicts on, and some type of performance metric is computed. The k estimates of performance are averaged to get an overall training error. We use repeated k-fold cross-validation, which repeats the procedure above more than once, randomly resampling the k subsets or folds each time. This method often leads to a reduction in variance for test set predictions, and it is feasible with the relatively small number of observations in our training data. We use 5-fold cross-validation with 5 repeats.

The performance metric we use for all model training and testing is the receiver operating characteristic (ROC) and associated area under the ROC curve (AUC). While model accuracy is a straightforward method for describing classifier performance, it is a weak measure when the classes are imbalanced because it treats all classes equally, meaning if 80% of the training data observations belong to one of two classes, then simply predicting the dominant class for all observations would give an accuracy of 80%. Accuracy also does not assess whether the classifier's errors are more often false positives or false negatives (Sokolova and Lapalme, 2009). For this reason, the performance criteria sensitivity and specificity are commonly used for assessing classification performance. Sensitivity (SNS) is a measure of the classifier's ability to correctly identify positive labels. It is a ratio of the number of true positives (TP) to the sum of true positives and false negatives (FN), as in

$$\text{SNS} = \frac{TP}{TP + FN}. \tag{6.12}$$

Specificity (SPC) is a measure of the classifiers ability to correctly identify negative labels, or the ratio between the number of true negatives (TN) and the sum of the TN and false positives (FP). It is given by

$$\mathrm{SPC} = \frac{TN}{TN + FP}. \tag{6.13}$$

The ROC curve is a convenient way to summarise the relationship between these two metrics to illustrate the trade-off between a classifier's true positive rate (TPR) and false positive rate (FPR), where the TPR is the same as the SNS, and the FPR can be written as $1 - \mathrm{SPC}$. These metrics are computed given the classifier's predicted probabilities for each class across a number of different classification thresholds, and the results are plotted to show how the TPR and FPR are related. The AUC is a scale-invariant and classification-threshold-invariant metric that can be used to compare the general performance of a classifier by computing the total area under the plotted ROC curve. The AUC ranges in value from 0 to 1, with values closer to 1 indicating perfect classification and values equal to 0.5 or below indicating poor classifier performance.

For the logistic regression model with elastic net regularisation, a sequence of 100 values for λ are evaluated to determine the overall complexity of the model (Hastie et al., 2015). A sequence of 10 α values are specified during the training procedure to find the best mix of ridge and lasso penalties. The hyperparameter to be tuned in the random forests model, the *mtry* parameter, is tuned across a sequence of values below and above $\sqrt{p}$ (Friedman et al., 2001). In both models, the hyperparameters are tuned such that the final selected model maximises the AUC.

After selecting final models, we evaluate their performance on the test set and compute ROC curves and AUC values for each.

6.4.4 Software

We conduct all of our analyses in the *R Statistics* environment (R Core Team, 2017). We use the following packages: for data processing and manipulation we use *dplyr* (Wickham et al., 2019), and for plotting we use *ggplot2* (Wickham, 2016). Our comparative cluster analysis is carried out using the *NbClust* package (Charrad et al., 2014). For our classification models, we use the *caret* (Kuhn, 2008) and *caretEnsemble* (Deane-Mayer and Knowles, 2016) packages, which incorporate both the *glmnet* package (Friedman et al., 2010) for regularisation and the *ranger* package (Wright and Ziegler, 2017) for random forests. Finally, we use the *pROC* package (Robin et al., 2019) to compute and measure classifier performance.

6.5 Results

This section first presents results of the cluster analysis, including the optimal number of clusters selected and visualisations of these. Second, it presents results of the classification analysis, including the predictive power of activity data and key activities associated with varying peak-period electricity demand patterns.

6.5.1 Cluster analysis results

Across all distance-algorithm combinations (including both k-means as well as hierarchical clustering with five different linkage methods), we evaluate results with several CVIs for 2 to 20 clusters. We find that of the CVIs with which we validate results, the gap statistic (Tibshirani et al., 2001) shows the least variability of index scores across the different distance-algorithm methods tested, suggesting it is robust to numerous clustering approaches. Index scores across all six CVIs are shown in the appendix (Figure 6.A.2).

The gap statistic is a statistical testing method that compares the change in within-cluster dispersion given different numbers of clusters k with that expected under a reference null distribution of the data. The optimal number of clusters is that which maximises the gap statistic (meaning the clustering structure is most different from a random uniform distribution).

We use the gap statistic to select the optimal number of clusters and assign final cluster membership. Figure 6.6 presents values of the gap statistic for $k = 2, \ldots, 20$. It shows that across the 13 distance-algorithm combinations, 2 clusters is returned as optimal in all cases.

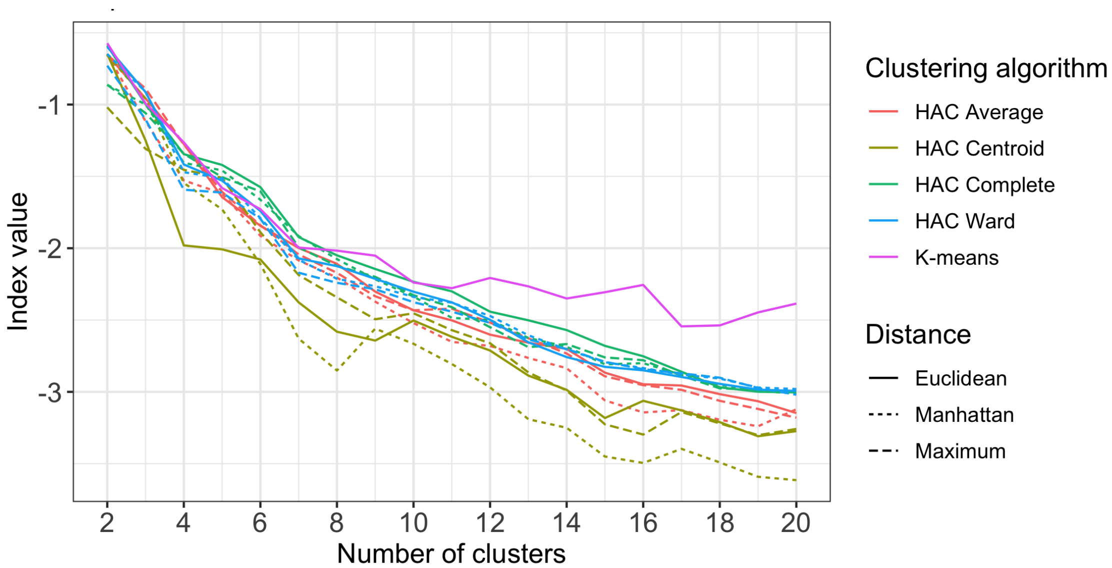

Figure 6.6: Gap statistic values across 13 distance-algorithm combinations, evaluated for 2 to 20 clusters.

We assign load profiles to the cluster (1 or 2) each is most frequently clustered in across these 13 methods. This approach avoids the need to select a specific distance-algorithm combination to assign final cluster membership, especially since the index performance across most of the methods for $k = 2$ is similar (see Figure 6.6).

Figure 6.7 plots all load profiles along with the mean value at each hour for the two clusters. From left to right, we present visualisations of each cluster and the cluster average for the normalised cumulative load profiles (the data on which the cluster analysis is performed), as well as the normalised clusters and the raw data. To highlight the methodological advantage of clustering cumulative rather than raw load profiles, we compute results for a counterfactual case, where the clustering is applied to de-minned profiles rather than cumulative de-minned profiles. This plot is included in the appendix (Figure 6.A.3) and shows how the cluster analysis is dominated by the magnitude of each profile's demand at 5 p.m. rather than its full shape throughout the evening period.

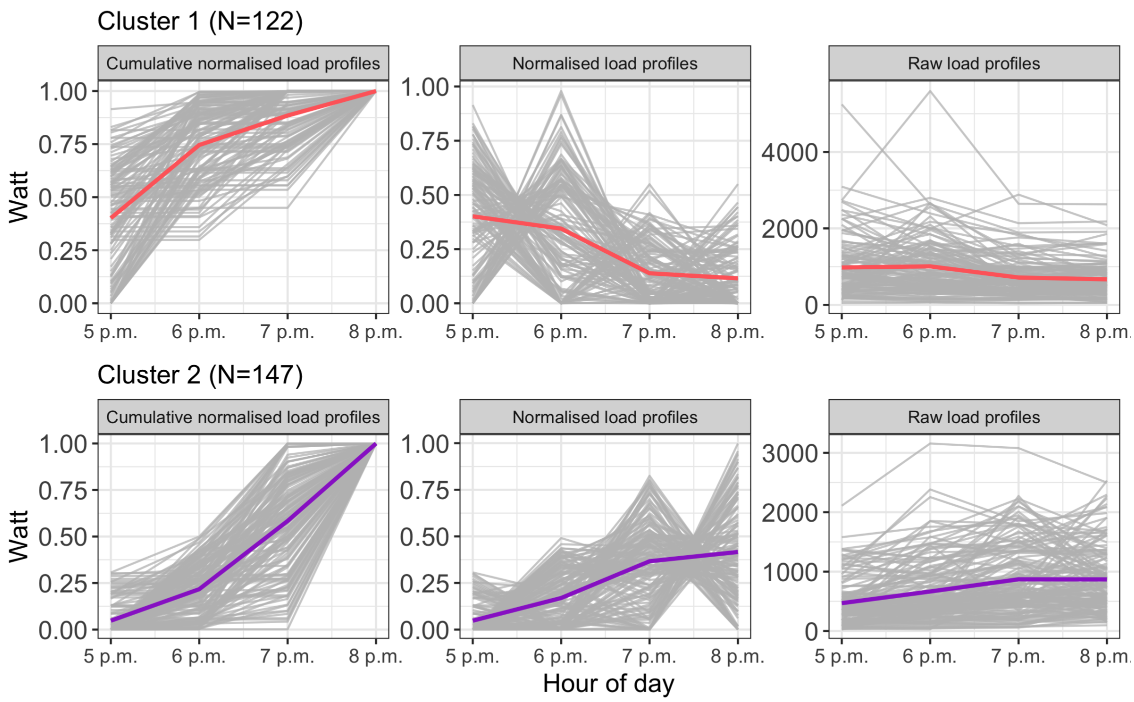

Figure 6.7: Cluster analysis results for $K = 2$, where plots show member load profiles and cluster averages for cumulative, normalized, and raw data.

As shown in the top panel of Figure 6.7, the first cluster consists of just under half the sample, and the average profile shape is characterised by higher discretionary demand in the earlier hours of the evening at 5 p.m. and then declining demand through the peak period. The second cluster is slightly larger in sample size, and its average profile shape is characterised by lower demand at 5 p.m., a steady increase through the peak period until 7 p.m., and stable demand thereafter.

In Figure 6.8, we show both cluster averages along with all load profiles on the same plot. Comparing the two, cluster 1's hourly demand of 978 W is nearly twice as high as cluster 2's (469 W) at 5 p.m. and around 1.5 times higher at 6 p.m. The cluster averages cross between 6 and 7 p.m., with Cluster 2's demand remaining around 150 W higher than

Cluster 1's from 7 to 9 p.m.

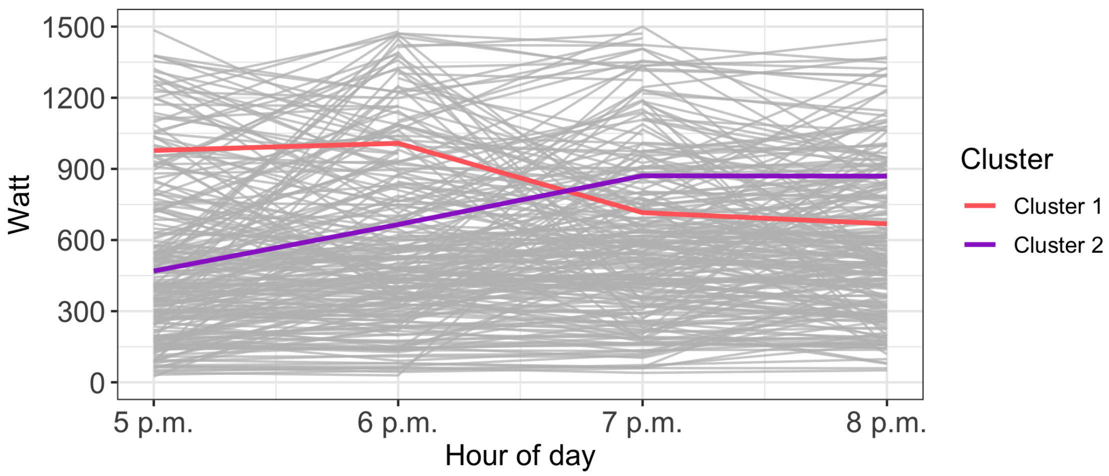

Figure 6.8: Raw load profiles plotted with cluster averages for $K = 2$.

The clustering approach demonstrated here fulfils the segmentation goal of identifying distinct clusters in terms of peak-period electric load profile shape. The results are especially useful for separating load profiles by the timing of peak demand, as cluster 1 should clearly be prioritised for interventions to reduce or shift demand during the early hours of the evening peak given its comparatively higher demand. Understanding the activity patterns driving these load profile groupings can help inform these efforts. For this reason, we next turn to the task of predicting cluster membership using occupant activity data.

6.5.2 Classification results

In both the logistic regression and the random forests models, we set up the prediction task as a binary classification problem, where the reference or base response category is membership in cluster 1, the early-peaking cluster. The models thus aim to predict the likelihood a load profile will be classified in the later-peaking cluster given the activities reported by occupants from 5–9 p.m.

In this section, we first present results related to model performance assessment and second on the predictive power of specific activities in the models.

Model performance assessment

We train our models on 70% of the data and use 5-fold repeated cross-validation to tune model hyperparameters and estimate test set error using AUC as the performance measure. The training results show that the best elastic net model is found with $\alpha = 0.22$ and $\lambda = 0.21$ and has an average cross-validation AUC of 0.73. The best random forests model is found with $mtry = 5$ and has an average cross-validation AUC of 0.72.

To assess the variability of these performance measures for the different samples taken during the cross-validation procedure, we use *Caret*'s built-in *resamples* function to compute the AUC metric across the 25 resamples (5 cross-validation folds with 5 repeats). Figure 6.9 presents a box-and-whisker plot of the variability in model training AUC for these 25 cross-validation samples.

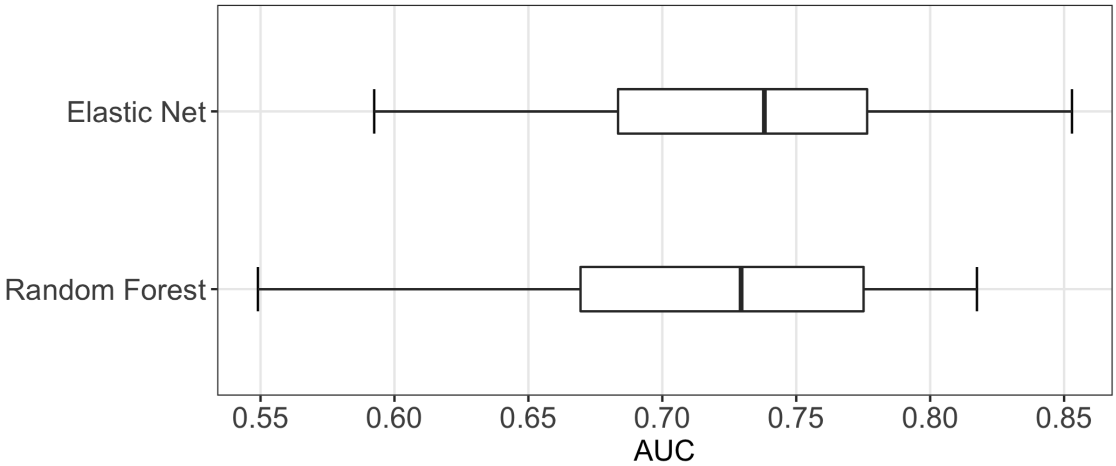

Figure 6.9: Box-and-whisker plot of training AUC across 25 cross-validation resamples.

Figure 6.9 shows that the elastic net model has a smaller interquartile range for training AUC than the random forests model, though there are not large performance differences between the models on the training data. Both do have relatively high overall ranges, however, suggesting the results are somewhat variable.

Next, we select the model with the tuning parameters that give the highest average cross-validated AUC in each case and use these to make predictions on the test set. We calculate ROC curves for each model's performance in predicting cluster membership for the test data. Figure 6.10 shows these curves and AUC values for each model. In both cases, the model performance declines slightly on the test data, which is typical in most model prediction contexts (James et al., 2013). The random forests model has a slightly higher test set AUC than the elastic net model, indicating that the elastic net model may have overfit the training data more than the random forests model.

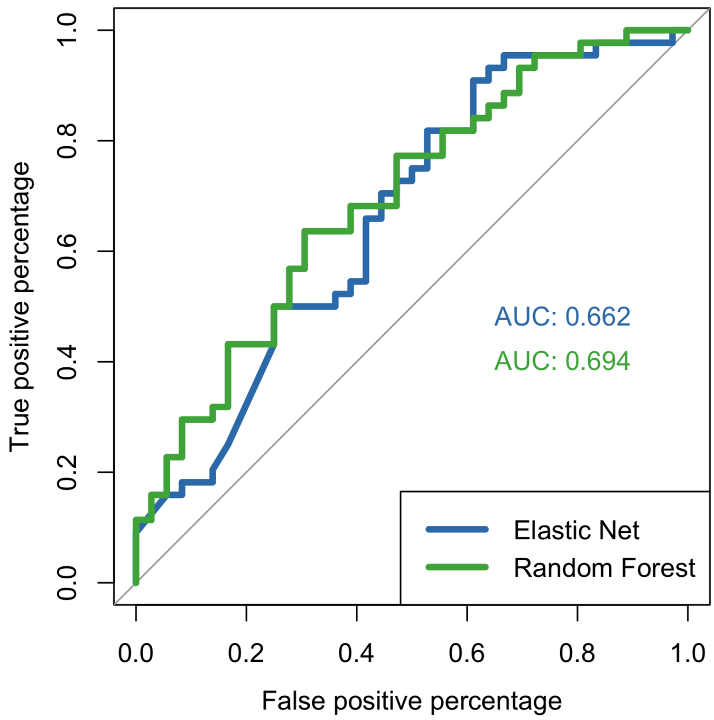

Figure 6.10: ROC curves and AUC values for elastic net and random forests performance on the test set.

Key activity variables

We inspect which activity predictors are influential in both classification models. For the elastic net model, we examine all non-zero coefficients that are selected, and for the random forests model, we assess variable importance using average impurity measured with the Gini index, as explained in Section 6.4.3. Table 6.3 presents predictors with non-zero coefficients in the elastic net model, including both their log($odds$) (β) and exponentiated coefficients (Exp[β]). These are ordered by hour and magnitude of β.

Hour (p.m.)	Activity group predictor	β	Exp(β)
5	Washing machine	−0.399	0.671
5	Socialising	−0.351	0.704
5	Laundry	−0.268	0.765
5	Eat	−0.151	0.860
5	Screen time	−0.131	0.877
5	Eating hot meal	−0.130	0.878
5	Tending domestic animals	0.107	1.113
5	Oven	−0.105	0.900
5	Cook on hob	−0.052	0.950
5	Other appliance	−0.036	0.964
5	Dishwasher	0.024	1.024
6	Eating hot meal	−0.337	0.714
6	Eat	−0.307	0.735
6	Got home	0.286	1.331
6	Other appliance	0.178	1.194
6	Pet	−0.162	0.850
6	Bath	−0.149	0.861
7	Oven	0.152	1.164
7	Childcare	−0.089	0.915
7	Screen time	−0.009	0.991
8	Laundry	0.388	1.474
8	Oven	0.281	1.325
8	Eating hot meal	0.012	1.012
	Intercept	0.499	1.648

Table 6.3: Predictors with non-zero coefficients in final logistic model with elastic net regularisation.

The elastic net regularisation technique delivers substantial sparsity for the logistic regression model, selecting 22 variables from the set of 177 candidate predictors. This result highlights that many of the hour-activity variables, perhaps owing to their low frequency, could be excluded from the model without sacrificing predictive performance. Elastic net regularisation is especially useful in this context, where $p \approx n$ after the testing data are held out. The technique helps balance the interests of including all hour-activity predictors as candidates for predicting peak-period electricity demand patterns while also selecting a more interpretable model.

Given that the base category in the model is membership in the earlier-peaking cluster, negative β coefficients signal a decrease in the likelihood that a load profile will belong to the later-peaking cluster or, similarly, an increased likelihood it will belong to the earlier-peaking cluster. For instance, the predictor '5-6 p.m. Eating hot meal' has $\beta = -0.13$ and $\text{Exp}[\beta] = 0.878$. This means that for each additional time this activity is reported in a given household, its likelihood of being classified in the later-peaking cluster decreases by 12%. In the case of the '5-6 p.m. Laundry' activity, each additional time it is reported decreases the likelihood a household will be classified in the later-peaking cluster by 23.5%.

As Table 6.3 shows, nearly all of the earlier (5-6 or 6-7 p.m.) selected predictors have negative coefficients, whereas most of the activity predictors reported later in the evening have positive coefficients. There are several exceptions, such as using the dishwasher from 5-6 p.m. or screen time from 7-8 p.m., but these generally have smaller coefficients.

In general, we find that meal-related activities are influential in the model, both in the case of earlier reports of meal activities showing strong associations with higher demand patterns in the early evening and later reports of meal activities showing strong associations with patterns in the later evening. Other key activities include laundry and washing machine use, screen time, and socialising.

Turning to the random forests model, we plot normalised variable importance in Figure 6.11, where the most important predictor is given a score of 1. We plot only the top 20 predictors for brevity.

As seen in the plot, many of the same variables that were influential in the elastic net model also have high variable importance scores in the random forests model. This is especially the case for the eating and meal-related predictors, screen time, and socialising.

An important difference between the linear elastic net model and the non-linear random forests model is that the predictors in the random forests model do not have coefficients and thus do not have a positive or negative sign and so cannot as easily be associated with increased probability of being in the early- or late-peaking cluster. Their importance is determined by how well they are able to split the predictor space. This lack of interpretability

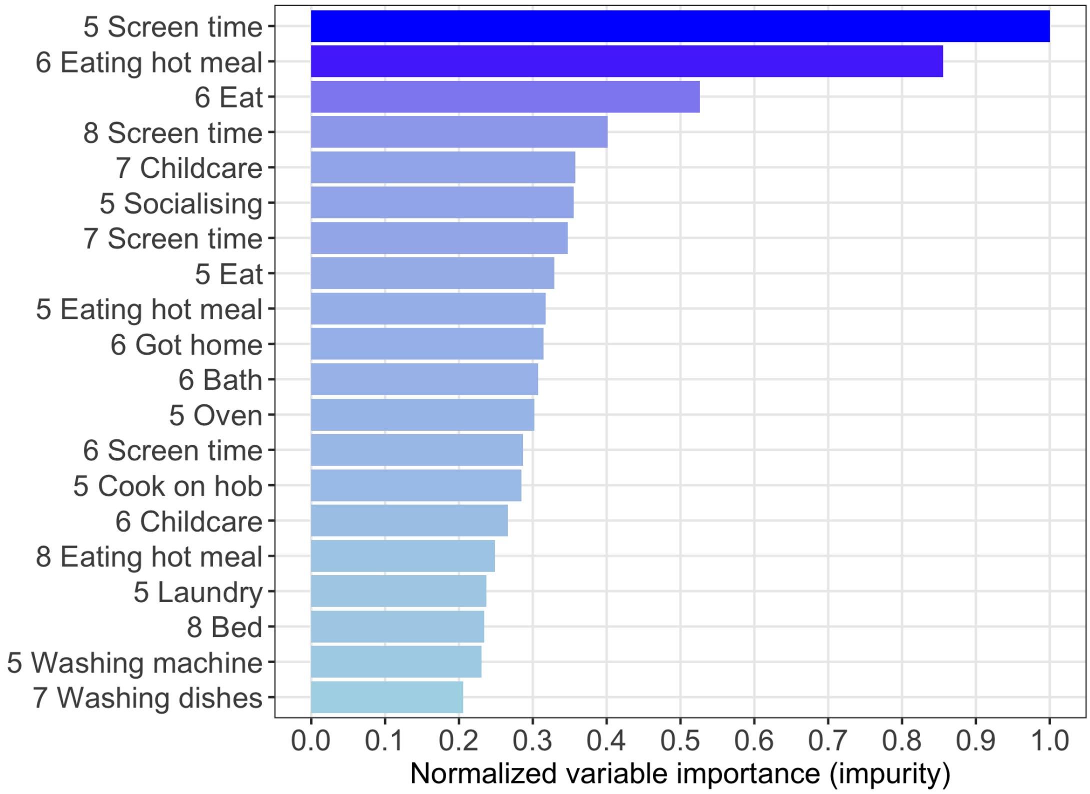

Figure 6.11: Normalised variable importance as determined by the Gini impurity method for the random forests model.

is one of the drawbacks of the random forests model, but considering many of the variables with high importance scores in this model are those selected in the elastic net model, it is reasonable to assume the direction of prediction is similar between the two. The random forests model, then, can provide further support to the findings on which activities are most influential for predicting peak-period electricity demand patterns.

6.6 Discussion and conclusions

This paper presents an integrated analysis to cluster a sample of 269 peak-period electric load profiles from homes in the U.K. We introduce a novel solution to capture temporal variations in demand by clustering cumulative load profiles. We show a new technique to appropriately use Euclidean distances with time-series data in order to cluster profiles based on their full shape during peak evening hours. This approach is beneficial for both its conceptual simplicity and its ease of implementation, and it does not appear to have been used previously for load profile segmentation.

We have shown how activity data can help to predict membership in early and late-peaking clusters. We introduce two classification methods that are appropriate for the predictive modelling task given analytical challenges inherent to this classification problem. The first is logistic regression with elastic net regularisation and the second is random forests. Both are used to evaluate the predictive power of activities on held out test data, as well as to identify key activity variables for predicting cluster membership.

The two classification algorithms used to predict cluster membership from occupant-reported activities show test set AUC values of 0.67–0.69. These scores are not indicative of highly accurate predictive models, but they also clearly outperform a null classification model, which would have an AUC of 0.5. This finding suggests that activity data as they are incorporated in this analysis show some promise for predicting patterns of electricity demand during peak hours. To the authors' knowledge, this paper presents the first attempt to show how detailed time-use data can be used in a cluster-and-classification approach to better understand the activities driving electricity demand patterns in the residential sector.

The identification of key activities driving evening demand is important given the complexity of the modelling task, for which simpler regression/classification techniques fail to give accurate and stable results. Both models confirm the importance of certain activities for predicting the shape of a household's evening weekday demand. The strongest link is found between food preparation and meal-time activities, especially cooking with energy-intensive appliances and eating hot meals. Other activities, including screen time, laundry, and socialising are shown to be important predictors, as well.

These empirical findings are useful for better understanding the potential for activity-led peak demand reduction and DSR in households. Usage curtailment potential varies considerably across customers, so understanding the factors that make customers more or less able to shift or curtail demand at different times of day can help to achieve a greater response. For instance, our results suggest that the households that contribute the most to U.K. system-wide peak demand are those that have earlier cooking and meal times. This knowledge is useful for considering appropriate interventions, as evening meal practices may be particularly challenging to shift. Design of targeted interventions can benefit from taking this knowledge into account.

Similarly, deeper insight into electricity demand patterns in homes and their relationship to occupant activities can inform the design of more suitable electricity tariffs. Assessing household responses to time-varying electricity tariffs is an active area of research, but there is still little agreement on how adoption of such tariffs influences electricity use patterns (Harding and Sexton, 2017). Previous analyses of a smaller sample of households from the METER study provide some evidence that electricity tariff choice is associated with different patterns of daily electricity use. Satre-Meloy et al. (2019) found that households enrolled in a renewable electricity tariff have lower average daily electricity demand, as well as lower demand during evening peak hours, than those on standard tariffs. The study also found that enrollment in an Economy 7 or 10 tariff, both of which charge lower prices for electricity use during nighttime hours, is linked to slightly higher nighttime electricity use. This evidence is based on a small number of households, however, as the share of those on either a renewable or Economy 7/10 tariff is less than 25%.

These initial findings are worth exploring in further detail with a larger sample of households enrolled in time-varying electricity tariffs. The implications of tariff selection on peak-period electricity usage patterns are especially important to understand given that time-varying tariffs are likely to become more common in the future with further deployment of smart metering infrastructure (Nicolson et al., 2018). Furthermore, as most of the research on household responses to dynamic electricity prices has focused on how much of a peak reduction is achievable, the data collected in this study can provide insight into what activities are shifted, curtailed, or substituted to actually deliver those reductions. For instance, understanding whether or not households on electricity tariffs that charge higher prices during peak hours are more likely to shift high-demand activities, such as laundry, cooking, or screen time activities—all of which were shown in this study to be predictive of evening peak-period demand patterns—can provide unique understanding on which types of activities may be more or less price-responsive. This knowledge could further be used to inform tariff design for different classes of customers, identify constraints to DSR in households, or enable better integration of demand-side resources in the electricity system (Flath et al., 2012).

Some limitations are present in the data and analysis. The sample size for this study is small, non-representative, and it consists of households that participate across different seasons and days of the week, excluding weekends (though most of our sample participated during the winter season). Day of week and seasonal differences are no doubt important for understanding the temporal aspects of activities and electricity demand, and future larger samples could be segmented to explore such variations. We also expect biases are present in our sample's activity reporting and in its socio-demographic makeup. Participant self-selection is a further source of bias.

Regarding the statistical methods we use, the predictive power demonstrated in the classification models is somewhat limited, though we expect this might improve with access to larger training data sets. Training AUC across cross-validation resamples also has high variability. Cluster analysis is typically applied to larger data sets, and for many of the CVIs used to evaluate our clustering results, the index scores are highly variable for different numbers of clusters. We use the Gap index for final cluster assignment because it shows less variability than the other CVIs, but our exploratory analysis of cluster results across different methods and CVIs shows that final segmentation of the load profiles is dependent on the clustering algorithm and distance metric chosen.

We expect the approach demonstrated in this paper to yield clearer insights when applied to a larger, more representative sample, and for this reason, data collection is ongoing. Future work should further assess the performance of both cluster and classification analyses presented in this paper. For the cluster analysis approach, we advocate for further comparisons of cumulative load profile clustering with other time-series-specific clustering approaches. For the classification analysis, we encourage similar analyses of the predictive power of time-use activity data. While we restrict our focus to predicting evening peak-period demand patterns, a similar approach could be applied to different times of day or to longer periods. The approach to aggregate and represent activities in the models should be tested in future analyses, especially the use of different time windows for aggregation (for instance, counting all activities reported every half hour rather than hour). Further work should also explore how results change if the modelling approach is set up as a regression rather than classification problem, especially if the key concern is predicting drivers of peak demand.

As one key objective of this analysis is to identify activities that could be specifically targeted in DSR programmes, future work should also aim to test whether and how much these kinds of targeted interventions can increase reductions in demand during peak hours. Collecting data on the kinds of activities that are shifted during DSR interventions provides numerous avenues for further research, such as exploring whether occupant responses to interventions are sustainable over longer time periods or understanding the unintended consequences that might result from these interventions. At present, utilities receive little feedback of this nature, and there thus remain substantial opportunities to better inform DSR interventions. We encourage more detailed investigations of high-resolution electricity data and occupant activities to ensure these deliver energy and cost saving benefits without adversely affecting vulnerable populations.

Data availability

The data used in this study are publicly available from the U.K. Data Service (Grünewald and Diakonova, 2019*b*).

Acknowledgements

The authors gratefully acknowledge funding for this research from the Engineering and Physical Sciences Research Council (EPSRC) [EP/M024652/1]. The authors declare no competing interests.

6.7 References

Albert, A. and Anderson, J. A. (1984), ‘On the existence of maximum likelihood estimates in logistic regression models’, *Biometrika* **71**(1), 1–10.

Albert, A. and Rajagopal, R. (2013), ‘Smart meter driven segmentation: What your consumption says about you’, *IEEE Transactions on Power Systems* **28**(4), 4019–4030.

Amasyali, K. and El-Gohary, N. M. (2018), ‘A review of data-driven building energy consumption prediction studies’, *Renewable and Sustainable Energy Reviews* **81**, 1192–1205.

Anderson, B. and Torriti, J. (2018), ‘Explaining shifts in U.K. electricity demand using time use data from 1974 to 2014’, *Energy Policy* **123**, 544–557.

Azaza, M. and Wallin, F. (2017), Smart meter data clustering using consumption indicators: Responsibility factor and consumption variability, *in* ‘Proceedings of the 9th International Conference on Applied Energy’, Vol. 142, Energy Procedia, Cardiff, United Kingdom, pp. 2236–2242.

Beckel, C., Sadamori, L., Staake, T. and Santini, S. (2014), ‘Revealing household characteristics from smart meter data’, *Energy* **78**, 397–410.

BEIS (2018), DUKES chapter 5: Statistics on electricity from generation through to sales, *in* ‘Digest of UK Energy Statistics (DUKES): Electricity’, Department for Business, Energy & Industrial Strategy, London, United Kingdom, pp. 111–153.

Biau, G. (2012), ‘Analysis of a Random Forests Model’, *Journal of Machine Learning Research* **13**, 33.

Caliński, T. and Harabasz, J. (1974), ‘A dendrite method for cluster analysis’, *Communications in Statistics* **3**(1), 1–27.

Cao, H.-A., Beckel, C. and Staake, T. (2013), Are domestic load profiles stable over time? An attempt to identify target households for demand side management campaigns, *in* 'IECON 2013 - 39th Annual Conference of the IEEE Industrial Electronics Society', Institute of Electrical and Electronics Engineers, Vienna, Austria, pp. 4733–4738.

Charrad, M., Ghazzali, N., Boiteau, V. and Niknafs, A. (2014), 'NbClust: An R package for determining the relevant number of clusters in a data set', *Journal of Statistical Software* **61**(6).

Chicco, G., Napoli, R. and Piglione, F. P. (2006), 'Comparisons among clustering techniques for electricity customer classification', *IEEE Transactions on Power Systems* **21**(2), 933–940.

Davies, D. L. and Bouldin, D. W. (1979), 'A cluster separation measure', *IEEE Transactions on Pattern Analysis and Machine Intelligence* **PAMI-1**(2), 224–227.

De Lauretis, S., Ghersi, F. and Cayla, J.-M. (2017), 'Energy consumption and activity patterns: An analysis extended to total time and energy use for French households', *Applied Energy* **206**, 634–648.

Deane-Mayer, Z. A. and Knowles, J. E. (2016), 'caretEnsemble: Ensembles of caret models'.

Dent, I., Aickelin, U. and Rodden, T. (2011), Application of a clustering framework to U.K domestic electricity data, *in* 'Proceedings of the 11th Annual Workshop on Computational Intelligence', UK Computational Intelligence, Manchester, United Kingdom, pp. 161–166.

Druckman, A., Buck, I., Hayward, B. and Jackson, T. (2012), 'Time, gender and carbon: A study of the carbon implications of British adults' use of time', *Ecological Economics* **84**, 153–163.

du Toit, J., Davimes, R., Mohamed, A., Patel, K. and Nye, J. M. (2016), 'Customer segmentation using unsupervised learning on daily energy load profiles', *Journal of Advances in Information Technology* **7**(2), 69–75.

Duda, R. O. and Hart, P. E. (1973), *Pattern Classification and Scene Analysis*, Wiley-Blackwell, New York.

Edwards, R. E., New, J. and Parker, L. E. (2012), 'Predicting future hourly residential electrical consumption: A machine learning case study', *Energy and Buildings* **49**, 591–603.

eurostat (2009), *Harmonised European Time Use Surveys: 2008 Guidelines*, Office for Official Publications of the European Communities, Luxembourg. OCLC: 698893637.

Farrar, D. E. and Glauber, R. R. (1967), 'Multicollinearity in regression analysis: The problem revisited', *The Review of Economics and Statistics* **49**(1), 92.

Figueiredo, V., Rodrigues, F., Vale, Z. and Gouveia, J. B. (2005), 'An electric energy consumer characterization framework based on data mining techniques', *IEEE Transactions on Power Systems* **20**(2), 596–602.

Flath, C., Nicolay, D., Conte, T., van Dinther, C. and Filipova-Neumann, L. (2012), 'Cluster analysis of smart metering data', *Business & Information Systems Engineering* **4**(1), 31–39.

Friedman, J., Hastie, T. and Tibshirani, R. (2001), *The Elements of Statistical Learning*, Vol. 1 of *Springer Series in Statistics*, Springer, Berlin, Germany.

Friedman, J., Hastie, T. and Tibshirani, R. (2010), 'Regularization paths for generalized linear models via coordinate descent', *Journal of Statistical Software* **33**(1), 1.

Gershuny, J. (2008), Time-use studies: Daily life and social change: Full research report, ESRC End of Award Report, RES-000-23-0704-A, ESRC, Swindon, United Kingdom.

Gouveia, J. P. and Seixas, J. (2016), 'Unraveling electricity consumption profiles in households through clusters: Combining smart meters and door-to-door surveys', *Energy and Buildings* **116**, 666–676.

Granell, R., Axon, C. J. and Wallom, D. C. H. (2015*a*), 'Clustering disaggregated load profiles using a Dirichlet process mixture model', *Energy Conversion and Management* **92**, 507–516.

Granell, R., Axon, C. J. and Wallom, D. C. H. (2015*b*), 'Impacts of raw data temporal resolution using selected clustering methods on residential electricity load profiles', *IEEE Transactions on Power Systems* **30**(6), 3217–3224.

Grunewald, P. and Diakonova, M. (2018*a*), 'The electricity footprint of household activities–implications for demand models', *Energy and Buildings* **174**, 635–641.

Grunewald, P. and Diakonova, M. (2018*b*), 'Flexibility, dynamism and diversity in energy supply and demand: A critical review', *Energy Research & Social Science* **38**, 58–66.

Grünewald, P. and Diakonova, M. (2019*a*), 'METER data', `http://www.energy-use.org/docs/data/`.

Grünewald, P. and Diakonova, M. (2019*b*), 'METER: U.K. household electricity and activity survey, 2016-2019: Secure access. [data collection].', *UK Data Service* **SN: 8475**. itemType: dataset.

Grünewald, P., Diakonova, M., Zilli, D., Bernard, J. and Matousek, A. (2017), What we do matters – a time-use app to capture energy relevant activities, *in* 'Proceedings of the 2017 ECEEE Summer Study on Energy Efficiency in Buildings', European Council for an Energy Efficient Economy, Presqu'île de Giens, France, pp. 2085–93.

Grünewald, P. and Layberry, R. (2015), Measuring the relationship between time-use and electricity consumption, *in* 'Proceedings of the 2015 ECEEE Summer Study on Energy Efficiency in Buildings', European Council for an Energy Efficient Economy, Presqu'île de Giens, France, pp. 2087–96.

Haben, S., Singleton, C. and Grindrod, P. (2016), 'Analysis and clustering of residential customers energy behavioral demand using smart meter data', *IEEE Transactions on Smart Grid* **7**(1), 136–144.

Harding, M. and Sexton, S. (2017), 'Household Response to Time-Varying Electricity Prices', *Annu. Rev. Resour. Econ.* **9**(1), 337–359.

Hartigan, J. A. (1975), *Clustering Algorithms*, Wiley Series in Probability and Mathematical Statistics, John Wiley & Sons, Inc., New York, NY.

Hastie, T., Tibshirani, R. and Wainwright, M. (2015), *Statistical Learning with Sparsity: The Lasso and Generalizations*, Chapman & Hall/CRC.

Hsu, D. (2015), 'Comparison of integrated clustering methods for accurate and stable prediction of building energy consumption data', *Applied Energy* **160**, 153–163.

Hubert, L. J. and Levin, J. R. (1976), 'A general statistical framework for assessing categorical clustering in free recall', *Psychological Bulletin* **83**(6), 1072–1080.

Iglesias, F. and Kastner, W. (2013), 'Analysis of similarity measures in times series clustering for the discovery of building energy patterns', *Energies* **6**(2), 579–597.

Jalas, M. and Juntunen, J. K. (2015), 'Energy intensive lifestyles: Time use, the activity patterns of consumers, and related energy demands in Finland', *Ecological Economics* **113**, 51–59.

James, G., Witten, D., Hastie, T. and Tibshirani, R. (2013), *An Introduction to Statistical Learning*, Vol. 103 of *Springer Texts in Statistics*, Springer New York, New York, NY.

Jin, L., Lee, D., Sim, A., Borgeson, S., Wu, K., Spurlock, C. A. and Todd, A. (2017), Comparison of clustering techniques for residential energy behavior using smart meter data, *in* 'Proceedings of the Thirty-First AAAI Conference on Artifical Intelligence', Association for the Advancement of Artificial Intelligence, San Francisco, CA.

Jin, L., Spurlock, A., Borgeson, S., Fredman, D., Hans, L., Patel, S. and Todd, A. (2016), Load shape clustering using residential smart meter data: A technical memorandum, Technical report, Lawrence Berkeley National Laboratory, Berkeley, CA.

Kuhn, M. (2008), 'Building Predictive Models in R Using the caret Package', *Journal of Statistical Software* **28**(1), 1–26.

Kwac, J., Flora, J. and Rajagopal, R. (2014), 'Household energy consumption segmentation using hourly data', *IEEE Transactions on Smart Grid* **5**(1), 420–430.

Kwac, J., Flora, J. and Rajagopal, R. (2018), 'Lifestyle segmentation based on energy consumption data', *IEEE Transactions on Smart Grid* **9**(4), 2409–2418.

Kwac, J., Tan, C., Sintov, N., Flora, J. and Rajagopal, R. (2013), Utility customer segmentation based on smart meter data: Empirical study, *in* 'Proceedings of the 2013 IEEE International Conference on Smart Grid Communications (SmartGridComm)', Institute of Electrical and Electronics Engineers, Vancouver, Canada, pp. 720–725.

Lee, J. D., Sun, D. L., Sun, Y. and Taylor, J. E. (2016), 'Exact post-selection inference, with application to the lasso', *The Annals of Statistics* **44**(3), 907–927.

Liao, T. W. (2005), 'Clustering of time series data—a survey', *Pattern Recognition* **38**(11), 1857–1874.

Liu, H., Yao, Z., Eklund, T. and Back, B. (2012), Electricity consumption time series profiling: A data mining application in energy industry, *in* P. Perner, ed., 'Advances in Data Mining. Applications and Theoretical Aspects', Springer, Berlin, Heidelberg, pp. 52–66.

Lockhart, R., Taylor, J., Tibshirani, R. J. and Tibshirani, R. (2014), 'A significance test for the lasso', *Annals of statistics* **42**(2), 413.

Ma, J. and Cheng, J. C. (2016), 'Identifying the influential features on the regional energy use intensity of residential buildings based on Random Forests', *Applied Energy* **183**, 193–201.

McKenna, E., Higginson, S., Grunewald, P. and Darby, S. J. (2017), 'Simulating residential demand response: Improving socio-technical assumptions in activity-based models of energy demand', *Energy Efficiency* .

McKenna, E. and Thomson, M. (2016), 'High-resolution stochastic integrated thermal–electrical domestic demand model', *Applied Energy* **165**, 445–461.

McLoughlin, F., Duffy, A. and Conlon, M. (2013), 'Evaluation of time series techniques to characterise domestic electricity demand', *Energy* **50**, 120–130.

McLoughlin, F., Duffy, A. and Conlon, M. (2015), 'A clustering approach to domestic electricity load profile characterisation using smart metering data', *Applied Energy* **141**, 190–199.

Milligan, G. W. and Cooper, M. C. (1985), 'An examination of procedures for determining the number of clusters in a data set', *Psychometrika* **50**(2), 159–179.

National Infrastructure Commission (2016), Smart Power, Technical report, London, United Kingdom.

Nicolson, M. L., Fell, M. J. and Huebner, G. M. (2018), 'Consumer demand for time of use electricity tariffs: A systematized review of the empirical evidence', *Renewable and Sustainable Energy Reviews* **97**, 276–289.

Ofgem (2010), Demand side reponse. A discussion paper, Technical report, Office of Gas and Electricity Markets, London, United Kingdom.

Palmer, J., Terry, N., Kane, T., Firth, S., Hughes, M., Pope, P., Young, J., Knight, D. and Godoy-Shimizu, D. (2013), Further analysis of the household electricity use survey - Electrical appliances at home: Tuning in to energy saving, Technical Report 475/09/2012, Cambridge Architectural Research Limited, elementenergy, and Loughborough University, London, United Kingdom.

Panapakidis, I. P., Papadopoulos, T. A., Christoforidis, G. C. and Papagiannis, G. K. (2014), 'Pattern recognition algorithms for electricity load curve analysis of buildings', *Energy and Buildings* **73**, 137–145.

Parti, M. and Parti, C. (1980), 'The Total and Appliance-Specific Conditional Demand for Electricity in the Household Sector', *The Bell Journal of Economics* **11**(1), 309–321.

Piao, M., Shon, H. S., Lee, J. Y. and Ryu, K. H. (2014), 'Subspace projection method based clustering analysis in load profiling', *IEEE Transactions on Power Systems* **29**(6), 2628–2635.

R Core Team (2017), 'R: A language and environment for statistical computing', R Foundation for Statistical Computing.

Raileanu, L. E. and Stoffel, K. (2004), 'Theoretical comparison between the Gini index and information gain criteria', *Annals of Mathematics and Artificial Intelligence* **41**(1), 77–93.

Räsänen, T., Voukantsis, D., Niska, H., Karatzas, K. and Kolehmainen, M. (2010), 'Data-based method for creating electricity use load profiles using large amount of customer-specific hourly measured electricity use data', *Applied Energy* **87**(11), 3538–3545.

Rhodes, J. D., Cole, W. J., Upshaw, C. R., Edgar, T. F. and Webber, M. E. (2014), 'Clustering analysis of residential electricity demand profiles', *Applied Energy* **135**, 461–471.

Richardson, I., Thomson, M., Infield, D. and Clifford, C. (2010), 'Domestic electricity use: A high-resolution energy demand model', *Energy and Buildings* **42**(10), 1878–1887.

Robin, X., Turck, N., Hainard, A., Tiberti, N., Lisacek, F., Sanchez, J.-C. and Müller, M. (2019), 'pROC: Display and analyze ROC curves'.

Rodrigues, F., Duarte, J., Figueiredo, V., Vale, Z. and Cordeiro, M. (2003), A comparative analysis of clustering algorithms applied to load profiling, *in* P. Perner and A. Rosenfeld, eds, 'Machine Learning and Data Mining in Pattern Recognition', Springer-Verlag, Berlin, Heidelberg, pp. 73–85.

Rousseeuw, P. J. (1987), 'Silhouettes: A graphical aid to the interpretation and validation of cluster analysis', *Journal of Computational and Applied Mathematics* **20**, 53–65.

Satre-Meloy, A. (2019), 'Investigating structural and occupant drivers of annual residential electricity consumption using regularization in regression models', *Energy* **174**, 148–168.

Satre-Meloy, A., Diakonova, M. and Grünewald, P. (2019), 'Daily life and demand: An analysis of intra-day variations in residential electricity consumption with time-use data', *Energy Efficiency* pp. 1–26.

Satre-Meloy, A. and Langevin, J. (2019), 'Assessing the time-sensitive impacts of energy efficiency and flexibility in the US building sector', *Environ. Res. Lett.* **14**(12), 124012.

Smith, B. A., Wong, J. and Rajagopal, R. (2012), A simple way to use interval data to segment residential customers for energy efficiency and demand response program targeting, *in* 'Proceedings of the 2012 ACEEE Summer Study on Energy Efficiency in Buildings', American Council for an Energy Efficient Economy, Pacific Grove, CA, p. 13.

Sokolova, M. and Lapalme, G. (2009), 'A systematic analysis of performance measures for classification tasks', *Information Processing & Management* **45**(4), 427–437.

Stephen, B., Mutanen, A. J., Galloway, S., Burt, G. and Jarventausta, P. (2014), 'Enhanced load profiling for residential network customers', *IEEE Transactions on Power Delivery* **29**(1), 88–96.

Swan, L. G. and Ugursal, V. I. (2009), 'Modeling of end-use energy consumption in the residential sector: A review of modeling techniques', *Renewable and Sustainable Energy Reviews* **13**(8), 1819–1835.

Taylor, J. and Tibshirani, R. J. (2015), 'Statistical learning and selective inference', *Proceedings of the National Academy of Sciences* **112**(25), 7629–7634.

Teeraratkul, T., O'Neill, D. and Lall, S. (2018), 'Shape-based approach to household electric load curve clustering and prediction', *IEEE Transactions on Smart Grid* **9**(5), 5196–5206.

Tibshirani, R., Walther, G. and Hastie, T. (2001), 'Estimating the number of clusters in a data set via the gap statistic', *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **63**(2), 411–423.

Torriti, J. (2014), 'A review of time use models of residential electricity demand', *Renewable and Sustainable Energy Reviews* **37**, 265–272.

Torriti, J., Hanna, R., Anderson, B., Yeboah, G. and Druckman, A. (2015), 'Peak residential electricity demand and social practices: Deriving flexibility and greenhouse gas intensities from time use and locational data', *Indoor and Built Environment* **24**(7), 891–912.

Tsekouras, G. J., Kotoulas, P. B., Tsirekis, C. D., Dialynas, E. N. and Hatziargyriou, N. D. (2008), 'A pattern recognition methodology for evaluation of load profiles and typical days of large electricity customers', *Electric Power Systems Research* **78**(9), 1494–1510.

Tso, G. K. and Yau, K. K. (2007), 'Predicting electricity energy consumption: A comparison of regression analysis, decision tree and neural networks', *Energy* **32**(9), 1761–1768.

U.S. Bureau of Labor Statistics (2018), American time use survey: Handbook of methods, Technical report, U.S. Department of Labor, Washington, D.C.

Viegas, J. L., Vieira, S. M., Melício, R., Mendes, V. and Sousa, J. M. (2016), 'Classification of new electricity customers based on surveys and smart metering data', *Energy* **107**, 804–817.

Wang, Y., Chen, Q., Kang, C., Zhang, M., Wang, K. and Zhao, Y. (2015), 'Load profiling and its application to demand response: A review', *Tsinghua Science and Technology* **20**(2), 117–129.

Wang, Z., Wang, Y., Zeng, R., Srinivasan, R. S. and Ahrentzen, S. (2018), 'Random forest based hourly building energy prediction', *Energy and Buildings* **171**, 11–25.

Wei, Y., Zhang, X., Shi, Y., Xia, L., Pan, S., Wu, J., Han, M. and Zhao, X. (2018), 'A review of data-driven approaches for prediction and classification of building energy consumption', *Renewable and Sustainable Energy Reviews* **82**, 1027–1047.

Wickham, H. (2016), *Ggplot2: Elegant Graphics for Data Analysis*, Springer-Verlag, New York.

Wickham, H., François, R., Henry, L. and Müller, K. (2019), 'Dplyr: A grammar of data manipulation'.

Widén, J., Lundh, M., Vassileva, I., Dahlquist, E., Ellegård, K. and Wäckelgård, E. (2009), 'Constructing load profiles for household electricity and hot water from time-use data—Modelling approach and validation', *Energy and Buildings* **41**(7), 753–768.

Widén, J. and Wäckelgård, E. (2010), 'A high-resolution stochastic model of domestic activity patterns and electricity demand', *Applied Energy* **87**(6), 1880–1892.

Wright, M. N. and Ziegler, A. (2017), 'Ranger: A fast implementation of random forests for high dimensional data in C++ and R', *Journal of Statistical Software* **77**(1), 1–17.

Xu, S., Barbour, E. and González, M. C. (2017), Household segmentation by load shape and daily consumption, *in* 'Proceedings of ACM SigKDD 2017 Conference', Halifax, Nova Scotia.

Yu, Z., Haghighat, F., Fung, B. C. M. and Yoshino, H. (2010), 'A decision tree method for building energy demand modeling', *Energy and Buildings* **42**(10), 1637–1646.

Zhang, T., Zhang, G., Lu, J., Feng, X. and Yang, W. (2012), 'A new index and classification approach for load pattern analysis of large electricity customers', *IEEE Transactions on Power Systems* **27**(1), 153–160.

Zhao, H.-X. and Magoulès, F. (2012), 'A review on the prediction of building energy consumption', *Renewable and Sustainable Energy Reviews* **16**(6), 3586–3592.

Zhou, K.-l., Yang, S.-l. and Shen, C. (2013), 'A review of electric load classification in smart grid environment', *Renewable and Sustainable Energy Reviews* **24**, 103–110.

Zou, H. and Hastie, T. (2005), 'Regularization and variable selection via the elastic net', *Journal of the Royal Statistical Society. Series B (Statistical Methodology)* **67**(2), 301–320.

6.A Supplemental figures

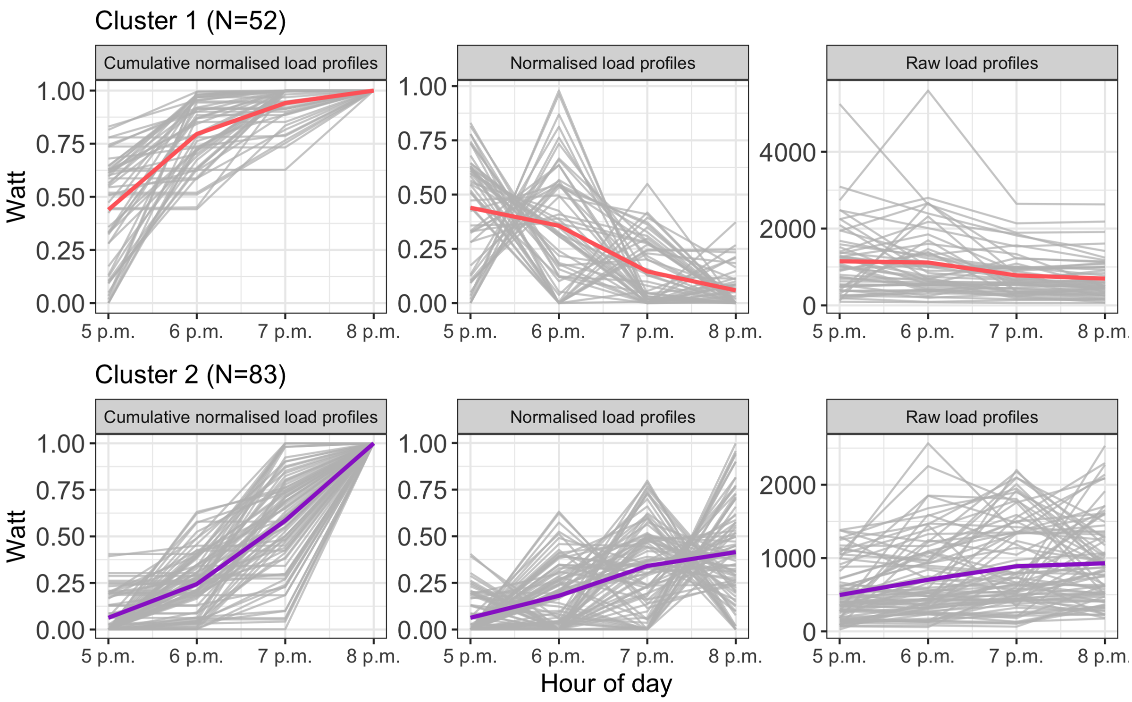

Figure 6.A.1: Clustering results using only load profiles from each household's first day of the study.

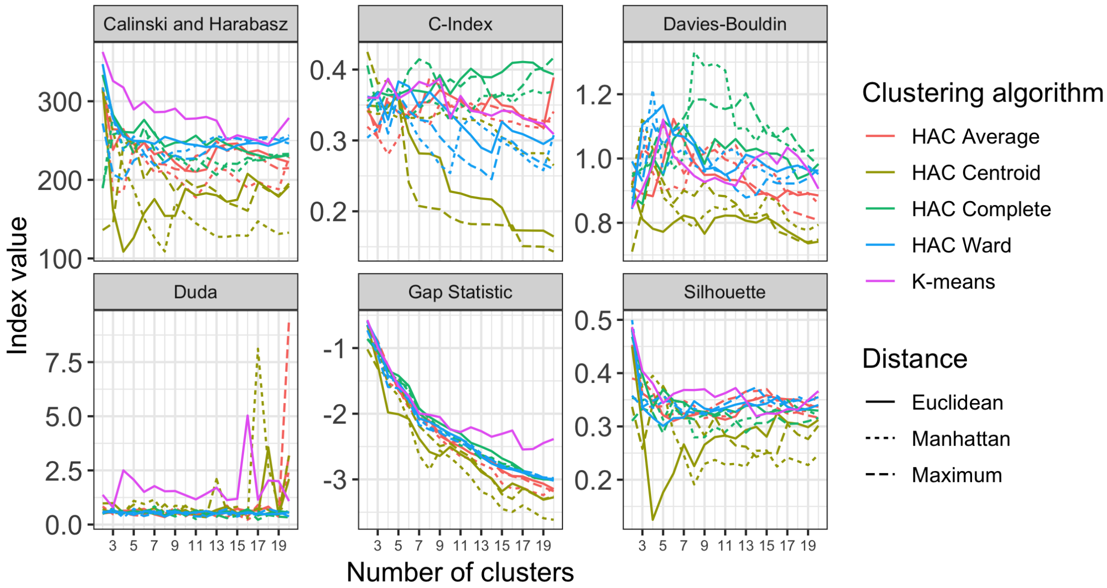

Figure 6.A.2: Results for cluster analysis evaluated with six CVIs for 2 to 20 clusters to show index score variability across different method combinations.

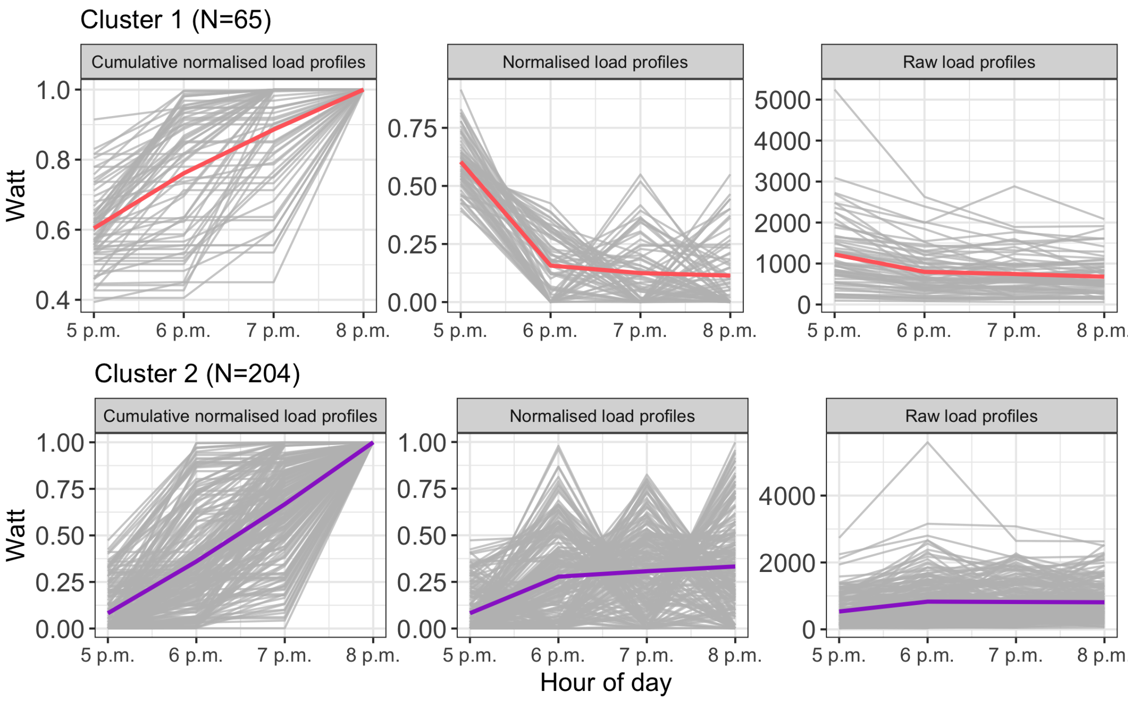

Figure 6.A.3: Results for cluster analysis applied to normalized load profiles (middle column plots) rather than cumulative normalized load profiles (left column plots).

6.B Supplemental tables

Index	Equation	Optimal number of clusters
Caliński and Harabasz (Caliński and Harabasz, 1974)	$\mathrm{CH}(q) = \frac{\mathrm{trace}(B_q)/(q-1)}{\mathrm{trace}(W_q)/(n-q)}$ where: • W_q is the within-group dispersion matrix for data clustered into q clusters, • B_q is the between-group dispersion matrix for data clustered into q clusters, • q is the number of clusters, • n is the number of observations	Maximum of index
C-index (Hubert and Levin, 1976)	$\mathrm{Cindex} = \frac{S_w - S_{\min}}{S_{\max} - S_{\min}}, S_{\min} \neq S_{\max},\ \mathrm{Cindex} \in (0, 1)$ where: • S_w is the sum of within-cluster distances $S_w = \sum_{k=1}^{q} \sum_{\substack{i,j \in C_k \\ i<j}} d(x_i, x_j)$ • S_{min} is the sum of the N_w smallest distances between all the pairs of points in the entire data set (there are N_t such pairs), • S_{max} is the sum of the N_w largest distances between all the pairs of points in the entire data set (there are N_t such pairs)	Minimum of index
Davies-Bouldin (Davies and Bouldin, 1979)	$\mathrm{DB}(q) = \frac{1}{q} \sum_{k=1}^{q} \max_{k \neq l}(\frac{\delta_k + \delta_l}{d_{kl}})$ where: • $k, l = 1, \ldots, q$ is the cluster number, • d_{kl} is the distance between centroids of clusters C_k and C_l, • δ_k is the dispersion measure of a cluster C_k	Minimum of index
Duda (Duda and Hart, 1973)	$\mathrm{Duda} = \frac{Je(2)}{Je(1)} = \frac{W_k + W_l}{W_m}$ where: • $J_e(2)$ is the sum of squared errors within clusters when the data are partitioned into two clusters, • $J_e(1)$ gives the squared errors when only one cluster is present. • It is assumed clusters C_k and C_l are merged to form C_m	Smallest number of clusters such that index > critical value (see Milligan and Cooper, 1985)
Gap statistic (Tibshirani et al., 2001)	$\mathrm{Gap}(q) = \frac{1}{B} \sum_{b=1}^{B} \log W_{qb} - \log W_q$ where: • B is the number of reference data sets generated using uniform prescription, • W_{qb} is the within-dispersion matrix defined as in Hartigan (1975)	Smallest number of clusters such that critical value ≥ 0
Silhouette (Rousseeuw, 1987)	$\mathrm{Silhouette} = \frac{\sum_{i=1}^{n} S(i)}{n},\ \mathrm{Silhouette} \in [-1, 1]$ where: • $S(i) = \frac{b(i) - a(i)}{\max\{a(i); b(i)\}}$ • $a(i)$ is average dissimilarity of the ith object to all other objects of cluster C_r • $b(i) = \min d_{iC_s}$ • d_{iC_s} is the average dissimilarity of the ith object to all objects of cluster C_s	Maximum of index

Table 6.B.1: Cluster validity indices (CVIs) for determining optimal number of clusters.

Chapter 7

Assessing the time-sensitive impacts of energy efficiency and flexibility in the U.S. building sector (Paper 4)

AVEN SATRE-MELOY, JARED LANGEVIN
PUBLISHED IN *Environmental Research Letters*

Abstract

The building sector consumes 75% of U.S. electricity, offering substantial energy, cost, and CO_2 emissions savings potential. New technologies enable buildings to flexibly manage electric loads across different times of day and season in support of a low-cost, low-carbon electric grid. Assessing the value of such technologies requires an understanding of building electric load variability at a higher temporal resolution than is demonstrated in previous studies of U.S. building efficiency potential. We adapt Scout, an open-access model of U.S. building energy use, to charactere sub-annual variations in baseline building electricity use, costs, and emissions at the national scale. We apply this baseline in time-sensitive analyses of the energy, cost, and CO_2 emissions savings potential of various degrees of energy efficiency and flexibility, finding that efficiency continues to have strong value in a time-sensitive assessment framework while the value of flexibility depends on assumed electricity rates, measure magnitude and duration, and the amount of savings already captured by efficiency.

7.1 Introduction

Residential and commercial buildings consumed 75% of U.S. electricity in 2018 (U.S. EIA, 2019) and are expected to drive nearly 70% of projected growth in U.S. electricity demand through 2040 (U.S. EIA, 2019). Buildings thus offer large potential savings in energy, operating costs, and CO_2 emissions for the U.S. electricity sector. Historically, national potential assessments of building efficiency have focused on estimating the annual energy savings delivered by policy instruments such as building codes and standards, equipment labelling, and technology research and development programs (Doris et al., 2009), with associated estimates of the economic value of energy efficiency reported on an annual basis, as well (Azevedo et al., 2013; EPRI, 2014, 2017; Gowrishankar and Levin, 2017; Granade et al., 2009; National Academy of Sciences et al., 2010; Wilson et al., 2017). Such annual potential assessments fundamentally assume that the operational impacts of efficient building technologies remain static across all hours of the day, days of the week, and seasons.

Advances in building technologies enable buildings to play a more active role in managing hourly electric loads to support a low-cost, modernised electric grid (FERC, 2018; Goldenberg et al., 2018). Smart controls and connectivity give buildings the ability to respond to grid signals and reduce or shift electricity consumption at certain times of day. These reductions and shifts may be achievable while providing comparable levels of core building services such as comfort to occupants, though service level impacts are highly dependent on building envelope and other factors such as occupant and operator preferences (Lund et al., 2015). In this paper, we refer to these smart, connected responses as the energy flexibility of a building, which the IEA defines as "the ability to manage its demand and generation according to local climate conditions, user needs and energy network requirements" (Jensen et al., 2017a, p.3). New forms of energy efficiency and flexibility technologies that provide grid services through load shedding and shifting may be an effective option to avoid electric system costs, such as capital costs for new power generation, operation and maintenance costs for existing generation, and capital costs for transmission and distribution upgrades (Dyson et al., 2015; O'Connell et al., 2014; Rahimi and Ipakchi, 2016). Energy flexible buildings may also support increased penetrations of renewable energy by reducing the risk of renewable power over-generation and curtailment, thus increasing its cost effectiveness (Denholm and Hand, 2011).

Because the cost of supplying electricity varies based on time of day and season (Mims et al., 2017), grid-focused valuations of building energy efficiency and flexibility must accordingly be assessed at a high temporal resolution. Existing national-scale energy modelling tools have limited ability to characterise the variations in building electricity use across sub-annual time intervals. For example, while the Electricity Market Module of the U.S. Energy Information Administration's (EIA) National Energy Modelling System (NEMS) (U.S. DOE, 2009) yields hourly estimates of total U.S. electricity demand, these estimates are not disaggregated to the building sector and its electric end uses, for which NEMS only yields annual estimates.

Moreover, the absence of a consistent framework for assessing the impacts of both energy efficiency and flexibility measures on baseline building electric loads makes it challenging to develop effective strategies for deploying these measures in tandem. Joint assessments of efficiency and flexibility measures are largely absent from the literature on national and regional energy demand, despite the need to understand potential trade-offs and synergies between the two approaches (Jensen et al., 2017a). Recent work places a strong focus on quantifying the flexible potential of buildings at an aggregate level without directly addressing the role of efficiency within the proposed methodologies (Jensen et al., 2017b; Junker et al., 2018; Lopes et al., 2016). Studies that do examine both energy efficiency and flexibility either rely on outdated baseline data sets and proprietary forecasts, place a limited focus on peak demand impacts (EPRI, 2009), or otherwise afford only qualitative descriptions of the relative impacts of efficiency and flexibility on electricity demand, their possible interactions, and related integration opportunities (Buckley, 2016; Goldman et al., 2010; Smith and Managan, 2012).

To address these gaps and limitations, we develop a new basis for quantitatively assessing the time-varying impacts of energy efficiency and flexibility on U.S. building energy use, energy costs, and CO_2 emissions. We map EIA projections of annual baseline building electricity use, cost, and emissions to a sub-annual basis, yielding estimates of hourly, seasonal, and regional variations in building electricity use that support time-sensitive valuation of energy efficiency and flexibility impacts at the national scale. We include an illustrative use of this updated baseline for the case of residential cooling to demonstrate how conventional energy efficiency measures compare to dynamic flexibility measures in terms of their electricity, cost, and emissions savings benefits in U.S. buildings.

To the authors' knowledge, this work is the first to develop a national baseline for time-sensitive valuation of energy efficiency and flexibility in the U.S. building sector. Our analysis framework and results can be used to demonstrate how next-generation building technologies that dynamically reshape energy loads across the day compare to traditional, static efficiency measures in terms of total energy, cost, and emissions savings potential. In this way, we aim to develop quantitative insights that can inform the emerging debate surrounding demand-side flexibility from buildings as part of energy policy making.

7.2 Methods

Hourly estimates of U.S. building energy use, operating costs, and CO_2 emissions are generated using Scout (scout.energy.gov), an open-source software programme[1] developed by the U.S. Department of Energy (U.S. DOE, 2019*b*). Scout estimates the national energy use, CO_2 emissions, and operating cost savings potential of emerging building energy conservation measures (ECMs) across a long time horizon (2015–2050); savings can be explored under multiple technology adoption cases nationally or for a subset of climate zones.

Given that Scout's analysis approach has been described in detail elsewhere (Langevin et al., 2019), we focus on describing the modifications we made to this approach to enable

[1]Scout's source code is publicly-available: `https://github.com/trynthink/scout`.

time-sensitive assessments (see Section 7.B in the appendix for an overview of Scout's analysis approach and baseline data). Specifically, we use typical daily energy load, price, and emissions shapes for each season and Scout climate region to re-apportion Scout's baseline annual energy, cost, and emissions totals, which reflect EIA Annual Energy Outlook Reference Case projections (U.S. EIA, 2018*c*), across all hours of a year.

Hourly energy load shapes are drawn from the Electric Power Research Institute (EPRI) End Use Load Shapes Library v5.0 (EPRI, 2012). *Average* hourly energy loads (kW) are normalised by annual electric demand across all hours of a certain day type in peak (May-September) and off-peak (October-April) seasons. Hourly energy loads are also broken out by pre-2004 North American Electric Reliability Corporation (NERC) region (U.S. DOE, 2017), facility type (residential or commercial), and energy end use (e.g., lighting, cooling, heating, etc.). A selection of the load shapes used is plotted in the appendix (Figure 7.C.1).

Hourly electricity price shapes correspond to active TOU rates in the U.S. Utility Rate Database (URDB) (NREL, 2018). Hourly TOU rates are broken out by customer type (residential, commercial, industrial), month of the year, day type (weekday, weekend), and U.S. Energy Information Administration (EIA) utility code.[2] A selection of the price shapes used is plotted in the appendix (Figure 7.C.2).

Finally, hourly marginal CO_2 emissions factors are drawn from a previous analysis of these factors for the U.S. electricity system (Siler-Evans et al., 2012). Emissions factors are broken out by post-2011 NERC region (U.S. DOE, 2017) and by three seasons, summer (May-September), winter (December-February), and intermediate (March-April, October-November). The full set of marginal emissions shapes used are plotted in the appendix (Figure 7.C.3).

To develop sub-annual estimates of U.S. building energy use, energy cost, and CO_2 emissions, these hourly energy load shapes, price shapes, and emissions factors are translated to a common temporal and spatial resolution and applied to Scout's default baseline data

[2]Only the residential and commercial TOU rates were used for this study.

for the sub-annual time segment(s) of interest.

First, a common set of three seasons is established across the data sets: summer (May-September), winter (December-February), and intermediate (March-April, October-November). These seasons match those used in the marginal emissions factor data, requiring modification to the EPRI and URDB data sets.[3]

Given these common seasonal definitions, hourly energy load, electricity price, and CO_2 emissions estimates are mapped to the annual timescale reflected in Scout's baseline data. For energy loads, we estimate the fraction of annual load, $\Phi^{\text{ld}}_{r,b,u,d,h}$, for each combination of pre-2004 NERC region r, building type b, end use u, season s, day type d, and hour h:[4]

$$\Phi^{\text{ld}}_{r,b,u,s,d,h} = \frac{L_{r,b,u,s,d,h}}{\sum_{s=1}^{S} \sum_{d=1}^{D} \gamma^{\text{ld}}_{s,d} \sum_{h=1}^{24} L_{r,b,u,s,d,h}} \tag{7.1}$$

where $L_{r,b,u,s,d,h}$ is the raw hourly load intensity from the EPRI database for pre-2004 NERC region r, building type b, end use u, season s, day type d, and hour h; D and S are the total sets of day types (weekday, weekend) and seasons (summer, winter, intermediate); and $\gamma^{\text{ld}}_{s,d}$ represents the total number of days per year that fall into day type d and season type s, further defined as:

$$\gamma^{\text{ld}}_{s,d} = 52.1429 N_{dw} \frac{N_{ds}}{365} \tag{7.2}$$

where the constants 52.1429 and 365 are the number of weeks and days per year, respectively, N_{dw} is the number of each day type per week (5 for weekdays, 2 for weekends), and N_{ds} is the number of days that fall under each season (153 in summer, 90 in winter, and 122 in the intermediate seasons).

For TOU electricity prices, we first assess the median and 5th/95th percentile price shapes for a given state st, building type b, season s, and day type d combination, where the

[3] In the case of the EPRI data set, it was assumed that the off-peak season (October-April) daily load shapes could be used to represent hourly load intensities for both the winter and intermediate seasons. In the case of the URDB data set, which breaks down electricity prices on a monthly basis, the average price shapes across all applicable months for each of the three seasons were used.

[4] For example: the fraction (between 0 and 1) of total annual residential cooling load in the MAIN region that occurs on weekdays between 1–2 p.m. across all summer months.

ratio of the maximum to minimum hourly electricity price was used to calculate percentiles for each combination. Electricity price shapes for each individual utility in the URDB are mapped to a U.S. state by finding the state that is associated with the utility's code in EIA form 861 (U.S. EIA, 2018*a*). Median and 5th/95th percentile electricity price shapes are then normalised by the average annual electricity price, yielding hourly electricity price intensities, $\Phi^{\mathrm{pr}}_{st,b,s,d,h}$, for a given combination of state st, building type b, season s, and day type d and hour h:[5]

$$\Phi^{\mathrm{pr}}_{st,b,s,d,h} = P_{st,b,s,d,h} \left(\frac{\sum_{s=1}^{S} \sum_{d=1}^{D} \gamma^{\mathrm{pr}}_{s,d} \sum_{h=1}^{24} P_{st,b,s,d,h}}{72} \right)^{-1} \tag{7.3}$$

where $P_{st,b,s,d,h}$ is the hourly electricity price from the URDB database for state st, building type b, season s, day type d, and hour h, and the parenthesised term represents the average electricity price across all 72 hours covered by our data for a given region and building type (3 seasons, 24 hours per season). $\gamma^{\mathrm{pr}}_{s,d}$ represents the total number of days per year that fall into day type d and season type s:

$$\gamma^{\mathrm{pr}}_{s,d} = \frac{N_{dw}}{7} N_{ms} \tag{7.4}$$

where the constant 7 is the number of days per week, N_{dw} is the number of each day type per week, and N_{ms} is the number of months in each season s.

For marginal emissions, hourly emissions factors in each season s are similarly normalised by the average marginal emissions factor across all hours and seasons, yielding hourly marginal emissions intensities, $\Phi^{\mathrm{mef}}_{rn,s,h}$, for a given combination of post-2011 NERC region rn, season s, and hour h:[6]

[5]For example: the electricity price on weekdays between 1–2 p.m. for the summer season in commercial buildings in Texas ($\Phi^{\mathrm{pr}}_{st,b,s,d,h} = 1$ represents an hourly price intensity that is equal to the average electricity price).

[6]For example: the marginal emissions intensity between 1–2 p.m. for the summer season in the ERCOT region ($\Phi^{\mathrm{mef}}_{rn,s,h} = 1$ represents an hourly marginal emissions intensity that is equal to the average emissions intensity).

$$\Phi^{\text{mef}}_{rn,s,h} = E_{rn,s,h} \left(\frac{\sum_{s=1}^{S} \sum_{h=1}^{24} E_{rn,s,h}}{72} \right)^{-1} \tag{7.5}$$

where $E_{rn,s,h}$ is the hourly marginal emissions factor (Siler-Evans et al., 2012) for post-2011 NERC region rn, season s, and hour h, and the parenthesised term represents the average marginal emissions factor across all 72 hours covered by our data for a given region.

The hourly annual load fractions ($\Phi^{\text{ld}}_{r,b,u,s,d,h}$) and hourly price and emissions intensities ($\Phi^{\text{pr}}_{st,b,s,d,h}$ and $\Phi^{\text{mef}}_{rn,s,h}$) calculated above are defined by different regional breakdowns that must be mapped to the American Institute of Architects (AIA) climate zone (z) breakdown used in Scout's baseline data. These climate zones, which are specified by number of cooling degree days and heating degree days, have been used historically by the U.S. Department of Energy for its Residential Energy Consumption Survey (RECS) U.S. DOE (2005). Pre-2004 NERC region (load shapes), state (price shapes), and post-2011 NERC region (marginal emissions factors) are mapped to AIA climate zone at the county resolution, enabling a population-weighted determination of the portion of each region that falls into each AIA climate zone.[7]

Finally, sub-annual baseline energy, emissions, and cost estimates are calculated by applying the hourly load fractions ($\Phi^{\text{ld}}_{z,b,u,s,d,h}$), price intensities ($\Phi^{\text{pr}}_{z,b,s,d,h}$), and emissions intensities ($\Phi^{\text{mef}}_{z,s,h}$) to Scout's annual baseline energy use, CO_2 emissions, and operating cost estimates for the electric fuel (f), as defined in eqs. 7.6–7.9:

$$E^{\text{base}}_{z,b,u,t,v,y,s,d,h} = E^{\text{base}}_{z,b,u,t,v,y} \Phi^{\text{ld}}_{z,b,u,s,d,h} \tag{7.6}$$

$$E^{\text{base-st}}_{z,b,u,t,v,y,s,d,h} = \frac{E^{\text{base}}_{z,b,u,t,v,y,s,d,h}}{SS_{f=\text{elec},y}} \tag{7.7}$$

$$C^{\text{base}}_{z,b,u,t,v,y,s,d,h} = E^{\text{base}}_{z,b,u,t,v,y,s,d,h} CI_{f,y} \Phi^{\text{mef}}_{z,s,h} \tag{7.8}$$

[7]The database used for this mapping is publicly available (Langevin and Harris, 2019) and draws from the U.S. Census Bureau (U.S. Census Bureau, 2017), the U.S. Environmental Protection Agency's eGRID database (U.S. EPA, 2018), and data from the Pacific Northwest National Laboratory used for mapping climate regions to counties (Baechler et al., 2015).

$$\psi^{\text{base}}_{z,b,u,t,v,y,s,d,h} = E^{\text{base}}_{z,b,u,t,v,y,s,d,h} FC_{b,y} \Phi^{\text{pr}}_{z,b,s,d,h} \tag{7.9}$$

where $E^{\text{base}}_{z,b,u,t,v,y,s,d,h}$, $C^{\text{base}}_{z,b,u,t,v,y,s,d,h}$, and $\psi^{\text{base}}_{z,b,u,t,v,y,s,d,h}$ are the total primary energy use, CO_2 emissions, and operating costs attributable to a given baseline stock segment (defined by climate zone z, building type b, end use u, technology t, and building vintage v) in projection year y, and to a given season s, day type d, and hour h within that projection year.[8] $SS_{f=\text{elec},y}$ is the site-to-source electricity factor in year y, which is used to translate primary energy use estimates to the site energy use estimates reported in our results, $E^{\text{base-st}}_{z,b,u,t,v,y,s,d,h}$. $CI_{f,y}$ and $FC_{b,y}$ are, respectively, the CO_2 emissions intensity for primary energy of baseline fuel type f in year y and the primary energy cost for building type b in year y. Calculations for site-to-source conversion factors, emissions intensities, and energy costs are further detailed in Section 7.B of the appendix.

To calculate the time-sensitive impacts of different energy efficiency and flexibility measures in Section 7.3.3, we define representative Scout ECMs that modify the sub-annual baselines based on percentage reductions to yield new estimates for eqs. 7.6–7.9. These modifications are specific to the type of measure applied (e.g., an efficiency measure applies a percentage load reduction evenly across all hours of the day, while a flexibility measure modifies the baselines based on percentage reductions during peak hours or percentage shifts from peak hours to off-peak hours). Given the variability of TOU rates in the URDB, we calculate savings potential for flexibility measures across a low, medium, and high savings potential, which are characterised by different measure duration and price intensities from the URDB. Further details are provided in Section 7.B of the appendix.

[8] For example: the total primary electricity use attributable to air source heat pump operation in single family homes in northern climates between 1–2 p.m. on summer weekdays in the year 2030.

7.3 Results

7.3.1 Hourly end-use electricity consumption, cost, and emissions totals

We first characterise hourly variations in electricity consumption by building end use in order to identify building loads that are substantial contributors to hourly variations in electricity demand, costs, and emissions. We focus on the residential sector and the year 2018 throughout the following analyses, but the model framework we use can make projections for both residential and commercial buildings for any year from 2015–2050. We include hourly end-use variations in electricity consumption for commercial buildings in 2018, as well as results for the year 2030 for both residential and commercial buildings, in the appendix (Figures 7.A.1–7.A.3).

Figure 7.1 presents hourly end-use electricity, cost, and emissions totals for residential buildings in 2018. Figure 7.1a shows that electricity demand peaks from 5–6 p.m., primarily driven by space cooling, which accounts for over 37% of the total load (excluding miscellaneous loads) during the peak hour. Minimum demand occurs from 3–4 a.m., when it is 1.7 times lower than at peak. Space heating and water heating are the largest contributors to demand during this hour. Across the day, thermal end uses show the largest temporal variations, whereas other end uses are comparatively flatter.

Figure 7.1b shows the total operating costs of electricity use across each hour, again disaggregated by building end use. This figure shows two cost values for each hour. The labelled totals above each stacked bar represent the cost of each hour's electricity demand under TOU pricing, where the total cost is derived using time-sensitive adjustment factors that weight hourly electricity costs based on an analysis of all existing residential TOU rates from the URDB. Here, we present cost totals using a rate shape that is the 50th percentile of all TOU rates in terms of peak to off-peak price ratio.

In addition to cost values based on the 50th percentile TOU rate structure, we also present hourly costs using the average residential retail rate for electricity, which was \$0.13/kWh in 2018 (U.S. EIA, 2018*d*). These costs are represented by a dashed line in

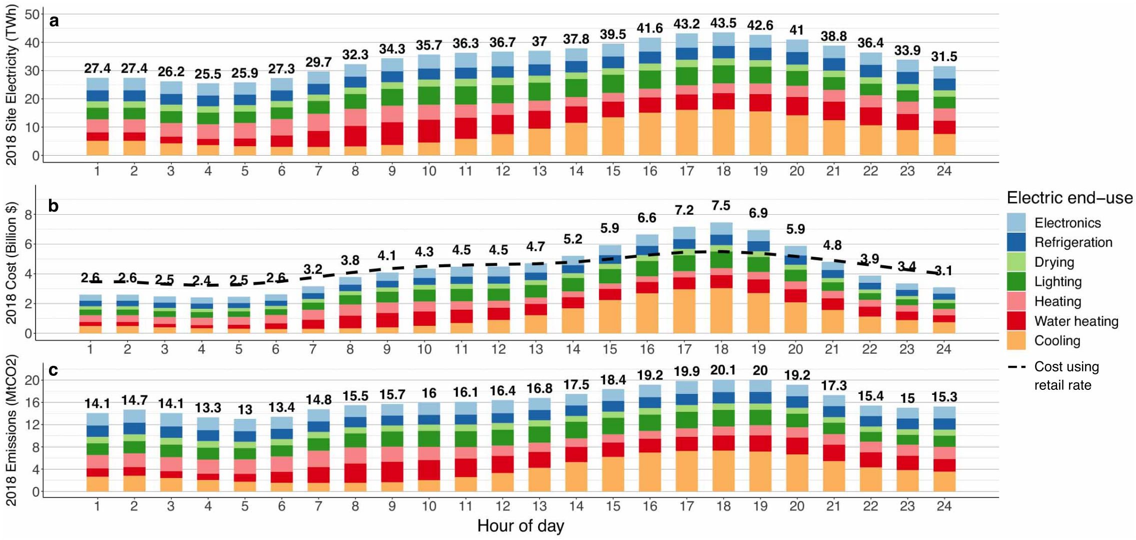

Figure 7.1: Hourly end-use electricity, cost, and emissions totals for residential buildings in 2018. (a), Total hourly site electricity consumption in terawatt hours (TWh). Each coloured bar represents a single end use, and bar labels indicate the total usage during each hour. (b), Total hourly operating costs in 2018 billion USD. End-use and labelled cost totals represent costs under full adoption of current TOU rates. Te dashed line represents cost totals using a national average retail rate for electricity. (c), Total hourly emissions from electricity generation in $MtCO_2$.

Figure 7.1b.

Estimating total costs under TOU pricing amplifies the peak to off-peak ratio seen in Figure 7.1a, as TOU rates increase during peak hours in order to reflect the higher costs to utilities of supplying electricity during these hours. The total cost during the peak hour is $7.5 billion, nearly three times higher than during 3–4 a.m. Space cooling again accounts for a substantial share of total costs, around 41% during the peak hour.

Comparing the cost totals under TOU pricing with costs totals under the average retail electricity price shows that full adoption of TOU rates nationally would increase costs by $8.3 billion between 2–8 p.m. This represents a 26% increase in electricity costs during the peak period and signals a significant increase in the potential cost savings for measures that reduce electricity use during these hours. Figure 7.1b reveals an opportunity to save electricity costs by shifting peak period electricity use to hours where costs under TOU pricing are below the dashed line (12–8 a.m.; 9–11 p.m.).

Total hourly emissions from electricity consumption are shown in Figure 7.1c. These totals are calculated by applying a marginal emissions scaling factor to each hour's electricity consumption total. The peak to off-peak ratio in marginal emissions, which is around 1.5, is lower than the ratios for either electricity use or costs. This is because marginal power sector emissions are, on average in the U.S. across regions and seasons, slightly higher during nighttime hours than during peak hours. This trend likely occurs because demand is low during nighttime and early-morning hours and coal is more often on the margin. When demand increases during morning and peak hours, gas-fired generators are more often on the margin because they have ramp rates that make them better suited to supplying this demand (Siler-Evans et al., 2012).

7.3.2 Seasonal and regional variations for electric space cooling

Given the large contributions of space cooling to hourly electricity consumption patterns in the residential sector, next we characterise how electric space cooling and its resulting costs vary seasonally and regionally across the U.S.

We include results for residential space heating in the appendix (Figure 7.A.4). We present results for seasonal and regional variations in space cooling and heating for commercial buildings there, as well (Figures 7.A.5–7.A.6).

For our analysis of seasonal variations, we split the year into summer (May-September) and winter/shoulder (October-April) seasons. Our regional analysis disaggregates national results to the five AIA climate zones as described in the Methods. The climate zones are ordered numerically from north to south.

As shown in Figure 7.2a, hourly summer electricity consumption for space cooling during the peak hour (5–6 p.m.) is five times larger than at its minimum (6–7 a.m.). Cooling electricity use during peak is concentrated in the southern and mid-Atlantic states, with CZ5 accounting for nearly 37% of demand and CZ3 and CZ4 accounting for an additional 44% during the peak hour.

The hourly cost variations shown in Figure 7.2b show an even larger peak to off-peak

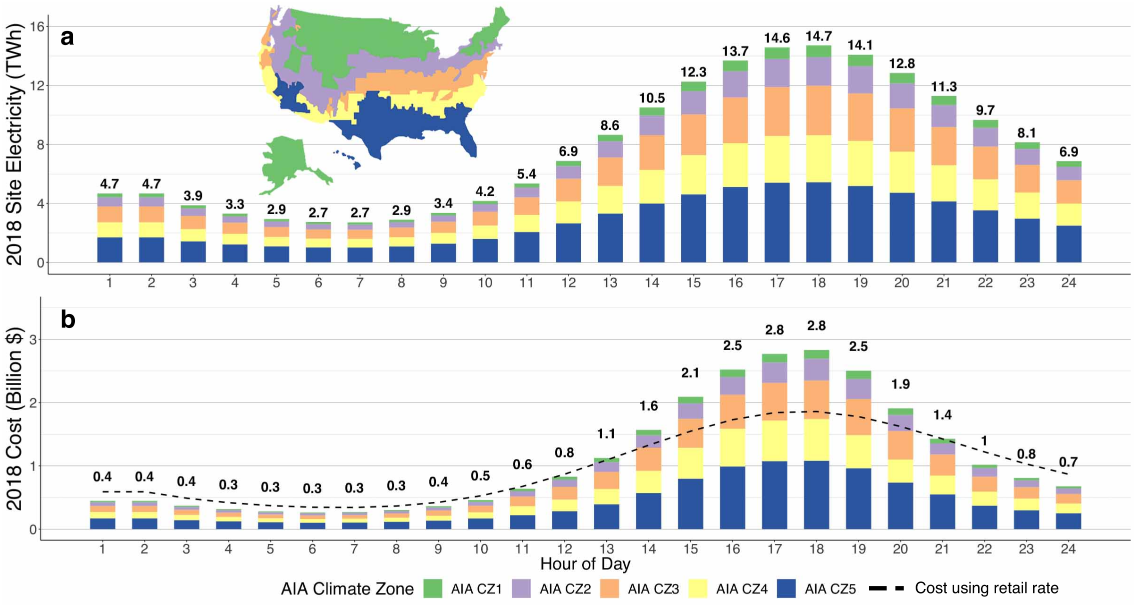

Figure 7.2: Hourly summer season (May-September) space cooling electricity and cost totals for residential buildings by climate zone in 2018. (a), Total hourly final electricity consumption in terawatt hours (TWh). Each colored bar represents one of the five U.S. climate zones, and bar labels indicate the total usage during each hour. (b), Total hourly operating costs in billion 2018 USD. Labeled cost totals represent costs under full adoption of TOU rates.

ratio for space cooling in summer (around 11 times higher during peak). This substantial ramp in hourly cooling costs is due to the alignment of high space cooling loads with the highest TOU prices for electricity during the early evening hours. The total costs for space cooling are again concentrated in climate zones 3–5.

7.3.3 Time-sensitive impacts of efficiency and flexibility measures

The hourly baselines of U.S. building electricity use developed in Figures 7.1–7.2 (and in Figures 7.A.1–7.A.6) can be used to assess the benefits of conventional efficiency technologies alongside new flexible building measures in terms of their electricity, cost, and emissions savings potential. We conduct two analyses to demonstrate how these baselines can be used to estimate demand, cost, and emissions impact potentials for residential space cooling measures. We choose residential cooling as an example focus because this end-use segment is both a substantial contributor to U.S. peak electricity demand as well as one that has clear efficiency and flexibility potential Chua et al. (2013), but the following analyses could be applied to other major end uses, seasons, and/or the commercial buildings sector.

In the first analysis, we show how national electricity demand and cost savings vary for hypothetical efficiency and flexibility measures with different magnitudes of impact on baseline residential building operations. We include similar estimates of savings potential for commercial cooling measures in 7.A (Figures 7.3.1–7.3.2).

In the second analysis, we demonstrate a more specific application of the sub-annual baselines presented in Sections 7.3.1–7.3.2 by considering example standalone efficiency and flexibility measures with building-level operational impacts that have been reported in previous studies. This analysis is not intended to definitively assess the savings potentials for each measure or for broader portfolios of efficiency and flexibility measures but rather presents an illustrative case of how such measures can be evaluated under a time-sensitive framework. It therefore relies on published point estimates of baseline load impacts, along with simplifying assumptions about building stock envelope efficiency levels, acceptable service thresholds, and occupant behaviour; we do not assess the uncertainty surrounding our estimated demand, cost, and emissions impacts given possible variations in these influencing factors.

Figure 7.3 illustrates the way the three types of measures we consider can reshape residential space cooling loads. Figure 7.3a shows a hypothetical static cooling efficiency measure, such as an efficient residential HVAC system, that reduces electricity use for space cooling evenly across all hours of the day. In Figure 7.3b, we represent a measure that sheds electricity use during peak hours (shown in the figure as 2–8 p.m.). Such a measure would entail, for instance, raising the thermostat cooling set point during this time window. Finally, Figure 7.3c shows a second flexibility measure that shifts electricity consumption from peak hours (2–8 p.m.) to the previous six hours, for example by pre-cooling the building. Note that the shift measure represented in Figure 7.3c assumes no net increase in daily energy load—e.g., reductions in peak load are directly translated into load increases of the same magnitude during the previous six hours.

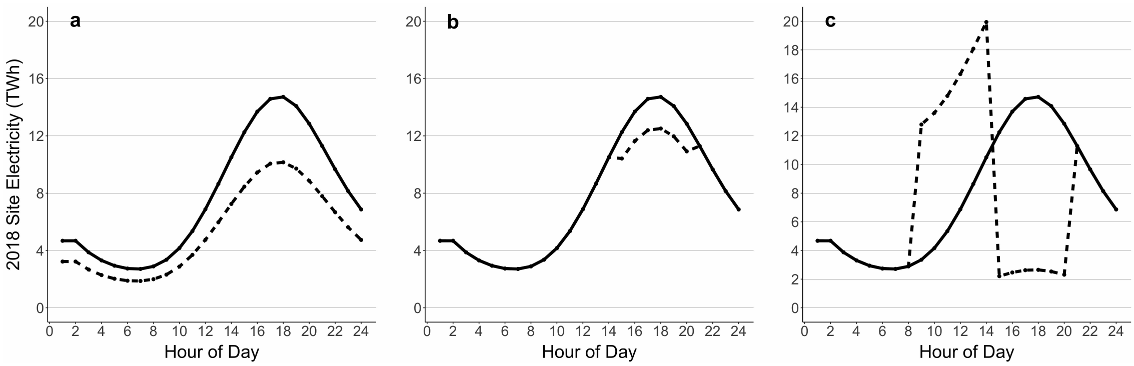

Figure 7.3: Load shapes representing hypothetical efficiency and flexibility measures. (a), static cooling efficiency measure that reduces electricity use evenly across hours. (b), dynamic cooling shed measure that reduces electricity use during peak hours (2–8 p.m.). (c), dynamic cooling shift measure that shifts electricity use from peak hours (2–8 p.m.) to the previous six hours. Note that this hypothetical measure does not assume an energy penalty.

We estimate the savings potential of these measures across varying assumptions about the magnitude of each measure's impact on baseline building operations and underlying TOU rate structures. In Figure 7.4, we present seasonal peak energy savings (TWh) and daily peak demand savings (GW) for the residential sector in summer across different magnitudes of peak reduction, corresponding to Figure 7.3b. Figure bars show total summer season savings in peak energy use while points indicate average daily peak demand savings, assuming a peak hour of 6 p.m. Total seasonal peak energy savings range from 8.2–41.1 TWh given a measure that sheds 10–50% of load during the hours of 2–8 p.m. Daily peak demand savings range from 9.2–46 GW, or around 1.2–5.9% of non-coincident peak summer demand in 2018 (U.S. EIA, 2017).

In Figure 7.5, we compare the cost savings potential of the measures represented in Figures 7.3a–7.3c. For the static efficiency measure (Figure 7.3a), we calculate cost savings for a 10%, 20%, and 30% reduction in baseline cooling load across all hours of the day. Both the shed (corresponding to Figure 7.3b) and shift (Figure 7.3c) measures are presented as savings ranges based on the magnitude of peak reduction, assumed duration of reduction,

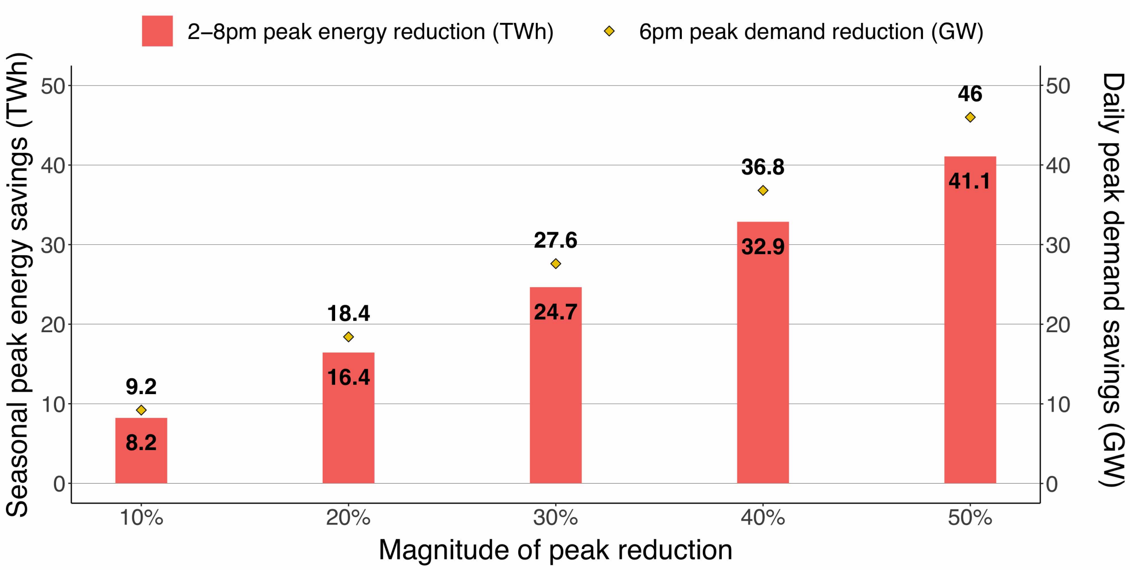

Figure 7.4: Summer season (May-September) peak energy and demand reduction potential for cooling in residential buildings at different magnitudes of peak energy use (TWh) and demand (GW) reduction. Bars show seasonal peak energy reduction potential during a peak period from 2–8 p.m. Points show daily peak demand reduction potential during a peak hour of 6 p.m.

and TOU rate structure. Consistent with Figure 7.3c, the shift estimates in Figure 7.5 reflect off-peak load increases that match on-peak load decreases—e.g., no additional daily energy penalty is assumed for load shifting. The influence of such penalties on total energy use outcomes is explored later in Table 7.1.

The cost savings for the static efficiency measures in Figure 7.5 total \$2.7, \$5.4, and \$8.2 billion for the 10%, 20%, and 30% load reductions, respectively. Achieving the cost savings of a 10% static efficiency measure using a dynamic measure instead would require a 15% peak reduction under the medium shed scenario (median rate structure; 2–8 p.m. peak period) or a 40% load shift under the high shift scenario (95th percentile rate structure; shift from 12–8 p.m. to 4 a.m.-12 p.m.). Similarly, the range in savings potential for a 25% shed measure is \$3.8–\$5.2 billion, which is the same as a static efficiency measure that saves 14–19% of electricity usage across all hours.

Regarding our analysis of several specific efficiency and flexibility measures, we present

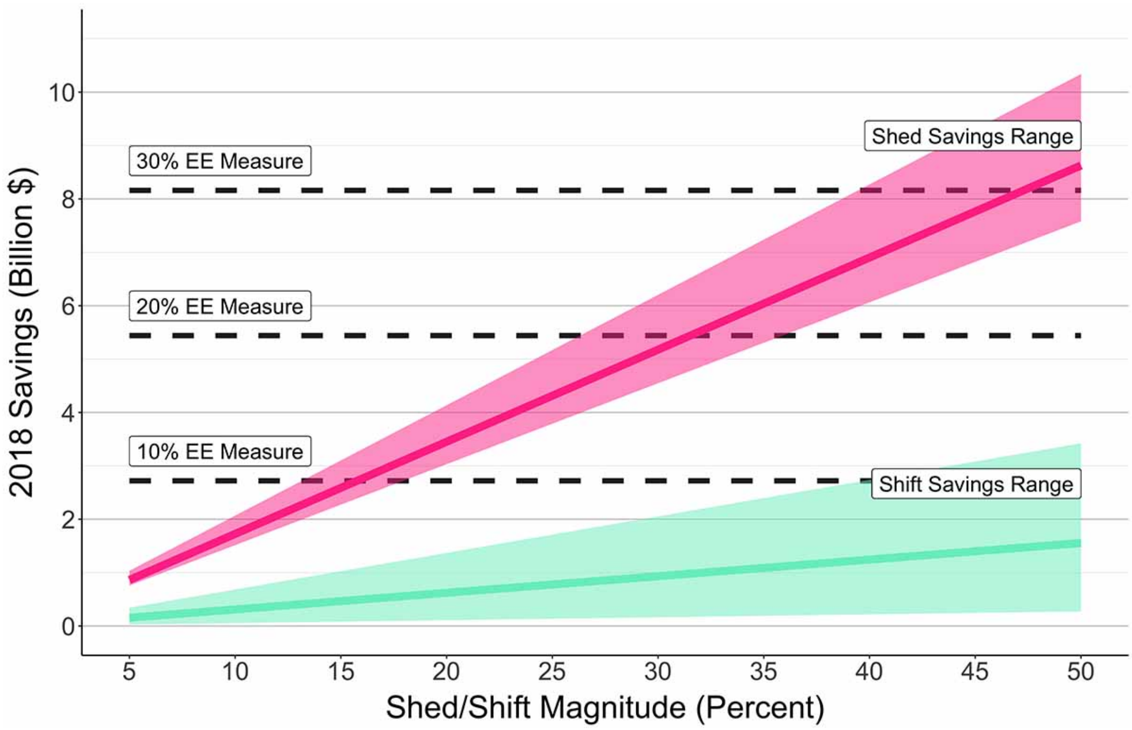

Figure 7.5: Technical cost savings potential for static efficiency and dynamic flexibility measures applied to summer season (May-September) residential cooling loads. Cost savings from static efficiency measures are shown with dashed lines representing three magnitudes of efficiency. The cost savings for both the shed and shift measures are presented as ranges for different magnitudes. The lower bound on the range is calculated using a four-hour shed/shift event and a flatter TOU rate structure (5th percentile of all URDB rates) while the upper bound is calculated using an eight-hour shed/shift event and a peakier TOU rate structure (95th percentile of all URDB rates). The solid lines show a six-hour shed/shift event using the median TOU rate structure.

seasonal electricity, cost, and emissions savings potentials along with daily peak demand reduction potentials for these measures in Table 7.1. The table includes the technology measures considered along with references for their estimated savings at the individual building level. It also includes assumptions related to measure magnitude and duration as well as the TOU rate structures used.

For the efficiency measure, we model a seasonal energy efficiency ratio (SEER) 18 cooling system, which can save up to 31% in cooling electricity consumption U.S. EIA (2018*g*). For the shed measure, we consider a thermostat set point adjustment, which could be either occupant-led or via direct load control (DLC) Chanana and Arora (2013). We assume a 15% reduction of cooling electricity usage during these hours Ivanov et al. (2013); Harding and Lamarche (2016). We base the shift measure on a previous study Turner et al. (2014), which simulates a mechanical pre-cooling measure in a home with thermal performance typical of new construction. We assume a 65% reduction in cooling load during peak hours for this measure and a peak-to-penalty energy ratio of 3.19, defined as the ratio of peak reduction to any *additional* increase in off-peak electricity demand beyond that attributable to pre-cooling. This off-peak load penalty is added on top of the pre-cooling electricity increase, which equals the magnitude of peak electricity reduction (see Section 7.B in the appendix for more details on how we translate this measure from the underlying study). Finally, we consider a combined efficiency and flexibility measure, which is based on an efficient cooling system with load shedding capabilities. For the more detailed assumptions on which these measures are based, especially those related to building envelope and service level impacts, we refer readers to the studies cited.

The results show that a static cooling efficiency measure delivers 57.3 TWh of site electricity savings, 29.8 GW of daily peak demand savings, \$8.1 billion in cost savings, and 26.4 $MtCO_2$ in emission savings. A flexible peak shedding measure, which reduces usage 15% from 4-8 p.m., yields 8.4 TWh, 14.4 GW, \$1.5 billion, and 3.8 $MtCO_2$ in savings. A flexible load shifting measure, which shifts 65% of the 4–8 p.m. load to the previous four hours via pre-cooling, leads to an overall increase in electricity use, costs, and CO_2 emissions but results in much larger daily peak demand savings (62.4 GW) than the other measures. A combined smart-controlled, high-efficiency HVAC system yields the largest electricity, cost, and emissions savings across these three metrics, totalling 63.1 TWh, \$9.2 billion, and 29.1 $MtCO_2$, though these savings are less than the sum of the efficiency and peak shedding measures when considered separately. This measure also delivers 39.8 GW of daily peak demand savings, the second highest magnitude of peak demand reduction after the pre-cooling shift measure.

Savings measure	Example technology	Assumptions	Estimated seasonal electricity savings (TWh)	Estimated daily peak demand savings (GW)	Estimated seasonal cost savings (Billion USD)	Estimated seasonal emissions savings ($MtCO_2$)
Static cooling efficiency improvement	SEER 18 cooling system (U.S. EIA, 2018*g*)	31% savings, all hours; median TOU rate structure	57.3†	29.8	\$8.1	26.4
Flexible cooling, peak shed	Thermostat set-point adjustment (Ivanov et al., 2013; Harding and Lamarche, 2016)	15% reduction, 4–8 p.m.; median TOU rate structure	8.4	14.4	\$1.5	3.8
Flexible cooling, peak shift	Thermal storage with pre-cooling (Turner et al., 2014)	65% shift from 4–8 p.m. to 4 hours earlier; 1.7°C thermostat turn-down; 3.19 peak-to-penalty ratio; median TOU rate structure	−11.4	62.4	−\$0.6	−5.5
Combined efficiency and flexibility measure	SEER 18 cooling system with smart-controlled load shedding capability	31% savings, all hours, plus 15% reduction, 4–8 p.m.; median TOU rate structure	63.1	39.8	\$9.2	29.1

† Reference totals for summer residential cooling in 2018: 184.8 TWh, 96.2 GW, \$26.2 bn USD, and 85.2 $MtCO_2$.

Table 7.1: Assessment of specific seasonal (May to September) cooling efficiency and flexibility measures under a time-sensitive framework to estimate their national electricity, cost, and CO_2 emissions savings potential in residential buildings in 2018.

7.4 Discussion

Examined across the summer season, the total electricity, cost, and emissions savings impacts of the peak shed and shift measures in Table 7.1 appear less favourable than those of installing more efficient cooling systems in residential buildings. Indeed, efficiency measures continue to have high value in a time-sensitive framework, as efficient cooling also delivers substantial reductions in peak demand. In the case of load shifting through pre-cooling, a slight increase in seasonal electricity use (around 6% of the total), costs, and emissions is observed, reflecting the influence of the off-peak energy penalty assumed for this measure. Nevertheless, this measure also yields the highest daily peak reductions, demonstrating how the benefits of these dynamic measures are variable under a time-sensitive framework depending on which time period is chosen for measure assessment. Moreover, Figure 7.5 shows that the cost implications of shed and shift measures depend heavily on the assumed magnitude and duration of peak savings for such measures, as well as assumed time-varying electricity rate structures. While the TOU rate shapes applied in this paper reflect current offerings from utilities, future rate structures may more heavily reward mid-day load increases (Seel et al., 2018), in which case pre-cooling measures would yield greater cost savings. Finally, while cost effectiveness was not the focus of the current analysis, measures that shed or shift loads may require lower incremental capital costs and thus deliver electricity savings more cost-effectively than conventional efficiency measures when paired with time-varying electricity rates (Dyson et al., 2015).

The results presented in Table 7.1 are primarily intended to show how the time-sensitive framework developed in this paper can be used to quantify the hourly and seasonal load, cost, and emissions impacts of specific efficiency and flexibility measures. As mentioned, we rely on point estimates of each measure's building-level operational impacts from previous studies, and a full uncertainty analysis of these estimates is beyond the scope of this paper. Nevertheless, we acknowledge the importance of undertaking more detailed analyses of measure operation at the building level, especially regarding the uncertainties inherent to measure performance in different building and climate contexts, as well as those related to a measure's expected service level impacts and the degree to which changes in building services would be accepted by building operators and occupants. Furthermore, we note that the full savings potentials estimated for the flexibility measures in Table 7.1 may not be realised when these measures are implemented in practice, as utilities will adjust the timing of load shifts and recovery periods to avoid creating new peaks in the system load shape.

Direct comparison of our findings with previous studies of U.S. building efficiency and flexibility impacts on hourly electric demand is precluded by the narrow focus in these studies on maximum peak demand reductions for a specific set of technology and programme deployment scenarios. For example, a study of national efficiency and demand response (DR) potential by EPRI (2009) estimates 218 GW of summer peak reduction from efficiency and DR measures by 2030, with 12.5 GW attributable to direct residential central air-conditioning control by utilities. A 2009 FERC study (2009) estimates 188 GW of summer peak reduction potential from DSR alone in 2019 but includes the industrial sector in this estimate. Another study of demand-side flexibility by the Rocky Mountain Institute (RMI) (Dyson et al., 2015) finds 49 GW of summer peak reduction potential from shedding 25% of residential air-conditioning loads in 2014. In all cases, the focus on maximum daily summer peak contrasts with our estimation of average daily summer peak; thus, the estimated impact on peak demand from a 25% residential cooling reduction in our study (about 25 GW, Figure 7.4) is smaller than that of the RMI study, for example. Moreover, these studies rely on data sets that are by now more than a decade old. Broadly speaking, however, the large potential these studies suggest for peak electric demand reductions from the building sector, and the prominent role of efficiency in these reductions in the EPRI study, are supported by our results.

The valuation approach presented in this paper is based on utility TOU rates for residential and commercial customers, assuming that such rates constitute a reasonable, readily-available proxy for the actual time-varying costs to utilities for supplying energy.[9] Future work could incorporate more direct proxies for the time-varying costs of electricity supply, such as wholesale locational marginal pricing (LMP), into our time-sensitive valuation framework. This research should consider how different valuation approaches might adjust the estimated savings potential of efficiency and flexibility measures, especially under high-renewable energy penetration scenarios, which are expected to reshape temporal variations in LMP considerably (Seel et al., 2018).

Similarly, the data used to apportion annual electricity end uses and emissions estimates to a sub-annual basis (see Section 7.2) will continue to be updated. Regarding sub-annual electric load data, reference models of whole building energy use from the U.S.

[9]Time-varying demand charges are also catalogued in the URDB and could similarly be used as a readily available utility cost proxy for time-sensitive efficiency valuation.

Department of Energy have recently been used to generate highly granular estimates of hourly energy demand across all major end uses, building types, and for all 8,760 hours of a typical meteorological year (U.S. DOE, 2013). A related effort at the U.S. national labs is generating hourly load profiles for thousands of prototypical buildings and calibrating these profiles against metered utility data. Regarding sub-annual emissions data, hourly marginal emissions intensities based on the U.S. Environmental Protection Agency's Continuous Emissions Monitoring System (CEMS) continue to be updated by the Carnegie Mellon Centre for Climate and Energy Decision Making and made publicly available via an online repository (Azevedo et al., 2017).

Future work will also focus on developing more specific definitions of energy flexibility measures, their impacts on building operations, and their cost effectiveness. While hourly load shape measurements for individual building technologies remain scarce, regionally-focused measurement efforts (e.g., NEEA, 2019) provide a basis for developing and validating measure load savings shapes, which can be extended to regions without readily available savings shape measurements using whole building energy simulation programs (U.S. DOE, 2019*a*). In developing such saving shapes, service level thresholds should be considered explicitly—at minimum by referring to relevant operational standards such as ASHRAE Standard 55 and 62 for thermal comfort and ventilation, respectively, and the IES Lighting Handbook for task illumination ANSI/ASHRAE (2016); ASHRAE (2017); DiLaura et al. (2011). Once developed and validated, measure savings shapes can be paired in Scout with associated data on technology capital costs and the latest estimates of time-varying electricity prices to explore the cost effectiveness of flexibility measures alongside that of conventional efficiency measures. Taken together, these updates to the modelling framework and its underlying data ensure that the current analysis is repeatable and that it remains relevant to policy questions concerning the role of building energy efficiency and flexibility in enabling a low-cost, low-carbon U.S. electricity system.

7.5 Conclusions

This paper develops hourly estimates of U.S. building electricity use, cost, and CO_2 emissions in order to facilitate the analysis of electricity saving measures with time-sensitive impacts alongside conventional, static efficiency measures. The national savings potentials for static energy efficiency measures and dynamic flexibility measures presented here enable a like-for-like comparison of next-generation building technologies that dynamically reshape energy loads with traditional efficiency technologies. Quantifying the magnitude of these measures' time-sensitive impacts across multiple metrics is critical to positioning the buildings sector as a key source of demand-side flexibility in a future that is likely to see increased stress on the power grid from climate change as well as higher penetrations of variable renewable energy generation (Denholm and Hand, 2011; Mukherjee and Nateghi, 2019).

Moreover, from a consumer perspective, a move towards electricity pricing that better reflects the real time-varying cost of energy supply will mean that reductions in electric load during peak demand hours deliver larger cost savings than reductions during off-peak hours, increasing the attractiveness of measures that yield impacts during high-price periods. An analysis framework that accounts for this time-varying value is essential for determining the cost effectiveness of large-scale energy efficiency or flexibility technology adoption in the building sector under the rate structures that are likely to emerge in a future U.S. electricity system that relies more heavily on distributed energy resources (AEE, 2018; Houghton et al., 2019).

Data availability

The data that support the findings of this study are openly available at DOI: `https://doi.org/10.5281/zenodo.3473478`.

Acknowledgements

This work was authored by The Regents of the University of California under Contract No. DE-AC02-05CH11231. Funding was provided by U.S. Department of Energy Office of Energy Efficiency and Renewable Energy Building Technologies Office. The views expressed in the article do not necessarily represent the views of the DOE or the U.S. Government.

7.6 References

AEE (2018), Rate design for a DER future: Designing rates to better integrate and value distributed energy resources, Technical report, Advanced Energy Economy, Washington, D.C.

ANSI/ASHRAE (2016), 'ASHRAE 62.1-2016, Ventilation for Acceptable Indoor Air Quality', `https://www.techstreet.com/ashrae/standards/ashrae-62-1-2016?product_id=1912838#document&ashrae_auth_token=`.

ASHRAE (2017), 'Standard 55-2017, Thermal Environmental Conditions for Human Occupancy', `https://www.ashrae.org/technical-resources/bookstore/standard-55-thermal-environmental-conditions-for-human-occupancy`.

Azevedo, I. L., Horner, N. C., Siler-Evans, K. and Vaishnav, P. T. (2017), Electriciy marginal factor estimates, Technical report, Carnegie Mellon University, Pittsburgh, PA.

Azevedo, I. L., Morgan, M. G., Palmer, K. and Lave, L. B. (2013), 'Reducing U.S. residential energy use and CO2 emissions: How much, how soon, and at what cost?', *Environmental Science & Technology* **47**(6), 2502–2511.

Baechler, M. C., Gilbride, T. L., Cole, P. C., Hefty, M. G. and Ruiz, K. (2015), High-performance home technologies: Guide to determining climate regions by county, Technical report, Pacific Northwest National Laboratory, Richland, WA.

Buckley, B. (2016), Putting more energy into peak savings: Integrating demand response and energy efficiency programs in the northeast and mid-atlantic, *in* 'Proceedings of the 2016 ACEEE Summer Study on Energy Efficiency in Buildings', American Council for an Energy Efficient Economy, Pacific Grove, CA, p. 13.

Chanana, S. and Arora, M. (2013), 'Demand response from residential air conditioning load using a programmable communication thermostat', *International Journal of Electrical and Computer Engineering* **7**(12), 7.

Chua, K., Chou, S., Yang, W. and Yan, J. (2013), 'Achieving better energy-efficient air conditioning – A review of technologies and strategies', *Applied Energy* **104**, 87–104.

Denholm, P. and Hand, M. (2011), 'Grid flexibility and storage required to achieve very high penetration of variable renewable electricity', *Energy Policy* **39**(3), 1817–1830.

DiLaura, D. L., Houser, K. W., Mistirck, R. G. and Steffy, G. R. (2011), *The Lighting Handbook, 10th Edition*, Illuminating Engineering Society, New York, NY.

Doris, E., Cochran, J. and Vorum, M. (2009), Energy efficiency policy in the United States: Overview of trends at different levels of government, Technical Report NREL/TP-6A2-46532, National Renewable Energy Laboratory, Golden, CO.

Dyson, M., Mandel, J., Lehrman, M., Bronski, P., Morris, J., Palazzi, T., Ramirez, S. and Touati, H. (2015), The Economics of demand flexibility: How "Flexiwatts" create quantifiable value for customers and the grid, Technical report, Rocky Mountain Institute, Basalt, CO.

EPRI (2009), Assessment of achievable potential from energy efficiency and demand response programs in the U.S. (2010 - 2030), Technical Report 1016987, Electric Power Research Institute, Palo Alto, CA.

EPRI (2012), 'End use load shapes', `http://loadshape.epri.com/enduse`.

EPRI (2014), U.S. energy efficiency potential through 2035, Technical Report 1025477, Electric Power Research Institute, Palo Alto, CA.

EPRI (2017), State level electric energy efficiency potential estimates, Technical Update 3002009988, Electric Power Research Institute, Palo Alto, CA.

FERC (2009), A national assessment of demand response potential, Staff Report, Federal Energy Regulatory Commission, Washington, D.C.

FERC (2018), 2018 assessment of demand response and advanced metering, Staff Report, Federal Energy Regulatory Commission, Washington, D.C.

Goldenberg, C., Dyson, M. and Masters, H. (2018), Demand flexibility: The key to enabling a low-cost, low-carbon grid, Technical report, Rocky Mountain Institute, Basalt, CO.

Goldman, C., Reid, M., Levy, R. and Silverstein, A. (2010), Coordination of energy efficiency and demand response, Technical Report LBNL-3044E, Lawrence Berkeley National Laboratory, Berkeley, CA.

Gowrishankar, V. and Levin, A. (2017), America's clean energy frontier: The pathway to a safer climate future, Technical Report 16-06-A, Natural Resources Defense Council, Washington, D.C.

Granade, H. C., Creyts, J., Derkach, A., Farese, P., Nyquist, S. and Ostrowski, K. (2009), Unlocking energy efficiency in the U.S. economy, Technical report, McKinsey Global Energy and Materials.

Harding, M. and Lamarche, C. (2016), 'Empowering consumers through data and smart technology: Experimental evidence on the consequences of time-of-use electricity pricing policies', *Journal of Policy Analysis and Management* **35**(4), 906–931.

Houghton, B., Salovaara, J. and Tai, H. (2019), Solving the rate puzzle: The future of electricity rate design, Technical report, McKinsey & Company.

Ivanov, C., Getachew, L., Fenrick, S. A. and Vittetoe, B. (2013), 'Enabling technologies and energy savings: The case of EnergyWise Smart Meter Pilot of Connexus Energy', *Utilities Policy* **26**, 76–84.

Jensen, S. Ø., Madsen, H., Lopes, R., Junker, R. G., Aelenei, D., Li, R., Metzger, S., Lindberg, K. B., Marszal, A. J., Kummert, M., Bayles, B., Mlecnik, E., Lollini, R. and Pasut, W. (2017a), *Annex 67: Energy Flexible Buildings*, International Energy Agency, Paris, France.

Jensen, S. Ø., Marszal-Pomianowska, A., Lollini, R., Pasut, W., Knotzer, A., Engelmann, P., Stafford, A. and Reynders, G. (2017b), 'IEA EBC Annex 67 energy flexible buildings', *Energy and Buildings* **155**, 25–34.

Junker, R. G., Azar, A. G., Lopes, R. A., Lindberg, K. B., Reynders, G., Relan, R. and Madsen, H. (2018), 'Characterizing the energy flexibility of buildings and districts', *Applied Energy* **225**, 175–182.

Langevin, J. and Harris, C. (2019), 'Scout geography mapping', `https://tinyurl.com/y2o7hzdo`.

Langevin, J., Harris, C. B. and Reyna, J. L. (2019), 'Assessing the Potential to Reduce U.S. Building CO2 Emissions 80% by 2050', *Joule* **3**(10), 2403–2424.

Lopes, R. A., Chambel, A., Neves, J., Aelenei, D. and Martins, J. (2016), 'A literature review of methodologies used to assess the energy flexibility of buildings', *Energy Procedia* **91**, 1053–1058.

Lund, P. D., Lindgren, J., Mikkola, J. and Salpakari, J. (2015), 'Review of energy system flexibility measures to enable high levels of variable renewable electricity', *Renewable and Sustainable Energy Reviews* **45**, 785–807.

Mims, N. A., Eckman, T. and Goldman, C. (2017), Time-varying value of electric energy efficiency, Technical Report 1398500, Lawrence Berkeley National Laboratory, Berkeley, CA.

Mukherjee, S. and Nateghi, R. (2019), 'A data-driven approach to assessing supply inadequacy risks due to climate-induced shifts in electricity demand', *Risk Analysis* **39**(3), 673–694.

National Academy of Sciences, National Academy of Engineering and National Research Council (2010), *Real Prospects for Energy Efficiency in the United States*, The National Academies Press, Washington, DC.

NEEA (2019), Residential building stock assessment II single-family homes report 2016-2017, Technical report, Northwest Energy Efficiency Alliance, Portland, OR.

NREL (2018), 'Utility Rate Database', `https://openei.org/wiki/Utility_Rate_Database`.

O'Connell, N., Pinson, P., Madsen, H. and O'Malley, M. (2014), 'Benefits and challenges of electrical demand response: A critical review', *Renewable and Sustainable Energy Reviews* **39**, 686–699.

Rahimi, F. and Ipakchi, A. (2016), 'Using a transactive energy framework: Providing grid services from smart buildings', *IEEE Electrification Magazine* **4**(4), 23–29.

Seel, J., Mills, A. D., Wiser, R. H., Deb, S., Asokkumar, A., Hassanzadeh, M. and Aarabali, A. (2018), Impacts of high variable renewable energy futures on wholesale electricity prices, and on electric-sector decision making, Technical Report 1437006, Lawrence Berkeley National Laboratory, Berkeley, CA.

Siler-Evans, K., Azevedo, I. L. and Morgan, M. G. (2012), 'Marginal emissions factors for the U.S. electricity system', *Environmental Science & Technology* **46**(9), 4742–4748.

Smith, K. and Managan, K. (2012), Energy efficiency and demand response: Working together in an integrated approach to managing energy, Technical report, Institute for Building Efficiency, Johnson Controls, Washington, D.C.

Turner, W., Walker, I. and Roux, J. (2014), Reducing residential peak electricity demand with mechanical pre-cooling of building thermal mass, Technical Report LBNL–180912, 1220527, Lawrence Berkeley National Laboratory, Berkeley, CA.

U.S. Census Bureau (2017), 'County population totals and components of change: 2010-2017', `https://www.census.gov/data/datasets/time-series/demo/popest/2010s-counties-total.html`.

U.S. DOE (2005), 'Maps', `https://www.eia.gov/consumption/residential/maps.php`.

U.S. DOE (2009), 'The National Energy Modeling System: An overview 2009', `https://www.eia.gov/outlooks/aeo/nems/overview/index.html`.

U.S. DOE (2013), 'Commercial and residential hourly load profiles for all TMY3 locations in the United States', `https://openei.org/doe-opendata/dataset/commercial-and-residential-hourly-load-profiles-for-all-tmy3-locations-in-the-united-states`.

U.S. DOE (2017), 'Form EIA-411 data', `https://www.eia.gov/electricity/data/eia411/`.

U.S. DOE (2019*a*), 'EnergyPlus', `https://energyplus.net/`.

U.S. DOE (2019*b*), 'Scout', `https://www.energy.gov/eere/buildings/scout`.

U.S. EIA (2017), 'Form EIA-411, "Coordinated bulk power supply program report," - Noncoincident peak load - Historic, actual and projected five-years', `https://www.eia.gov/electricity/data/eia411/pdf/peak_load_2016.pdf`.

U.S. EIA (2018*a*), 'Annual electric power industry report, form EIA-861 detailed data files', `https://www.eia.gov/electricity/data/eia861/`.

U.S. EIA (2018*b*), 'Annual Energy Outlook 2018: Table 18 - Energy-related carbon dioxide emissions by sector and source', `https://www.eia.gov/outlooks/aeo/data/browser/#/?id=17-AEO2019&cases=ref2019&sourcekey=0`.

U.S. EIA (2018*c*), 'Annual Energy Outlook 2018: Table 2 - Energy consumption by sector and source', https://www.eia.gov/outlooks/aeo/data/browser//#/?id=2-AEO2018&sourcekey=0.

U.S. EIA (2018*d*), 'Annual Energy Outlook 2018: Table 3 - Energy prices by sector and source', https://www.eia.gov/outlooks/aeo/data/browser/#/?id=3-AEO2018&cases=ref2018&sourcekey=0.

U.S. EIA (2018*e*), 'Commercial demand module of the national energy modeling system: Model documentation', https://www.eia.gov/outlooks/aeo/nems/documentation/commercial/pdf/m066(2017).pdf.

U.S. EIA (2018*f*), 'Residential demand module of the national energy modeling system: Model documentation 2018', http://www.osti.gov/servlets/purl/563838-kJEiRz/webviewable/.

U.S. EIA (2018*g*), 'Updated buildings sector appliance and equipment costs and efficiency', https://www.eia.gov/analysis/studies/buildings/equipcosts/pdf/full.pdf.

U.S. EIA (2019), Annual Energy Outlook 2019 with projections to 2050, Technical report, U.S. Energy Information Administration, Washington, D.C.

U.S. EPA (2018), 'eGRID subregion representational map', https://www.epa.gov/energy/egrid-subregion-representational-map.

Wilson, E., Christensen, C., Horowitz, S., Robertson, J. and Maguire, J. (2017), Electric end-use energy efficiency potential in the U.S. single-family housing stock, Technical Report NREL/TP–5500-65667, 1339938, National Renewable Energy Laboratory, Golden, CO.

7.A Supplemental figures

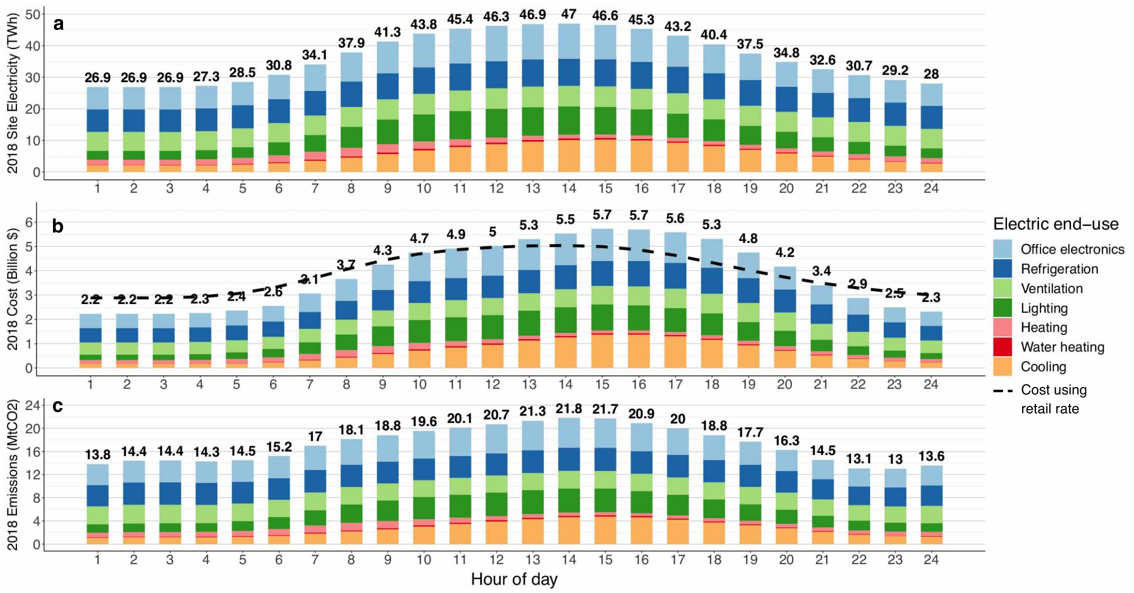

Figure 7.A.1: Hourly end-use electricity, cost, and emissions totals for commercial buildings in 2018. Figure elements and interpretation are the same as in Figure 7.1.

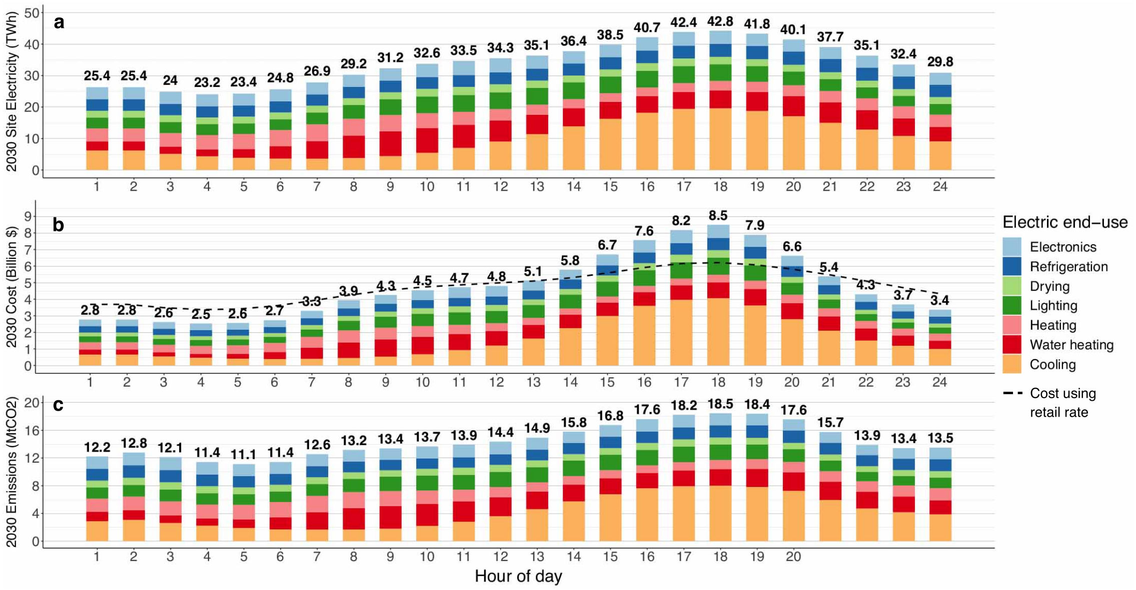

Figure 7.A.2: Hourly end-use electricity, cost, and emissions totals for residential buildings in 2030. Figure elements and interpretation are the same as in Figure 7.1.

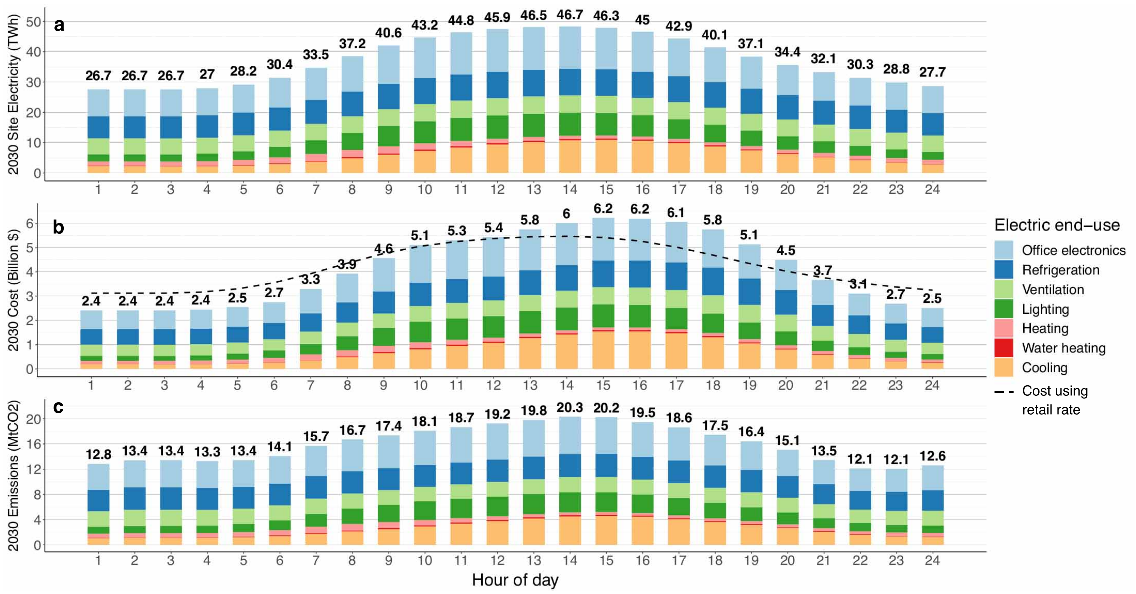

Figure 7.A.3: Hourly end-use electricity, cost, and emissions totals for commercial buildings in 2030. Figure elements and interpretation are the same as in Figure 7.1.

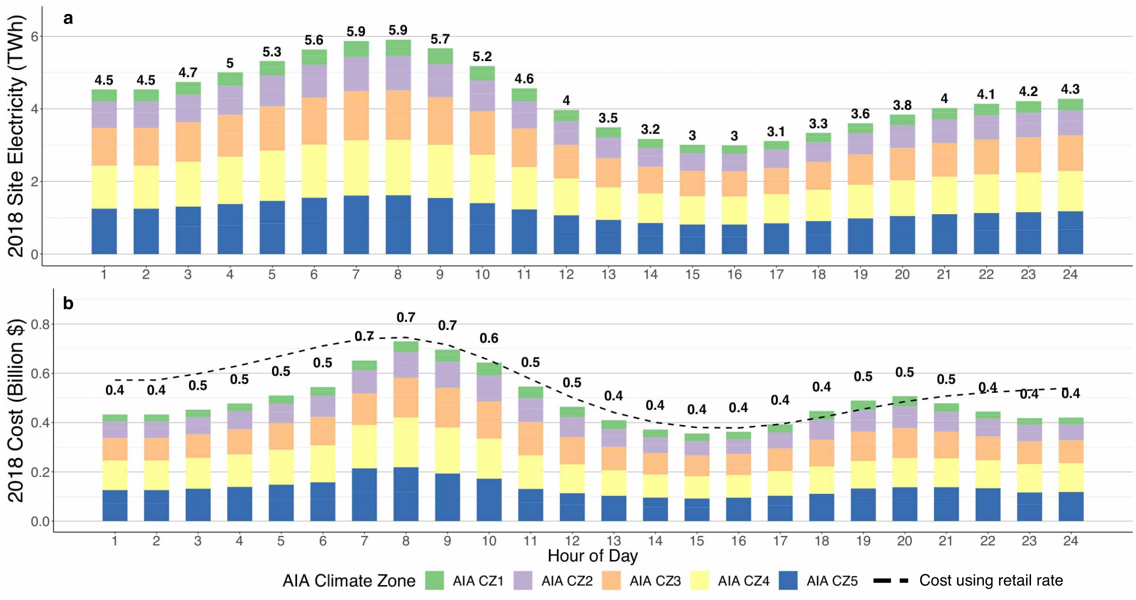

Figure 7.A.4: Hourly winter/intermediate season (October-April) space heating electricity and cost totals for residential buildings by climate zone in 2018. Figure elements and interpretation are the same as in Figure 7.2.

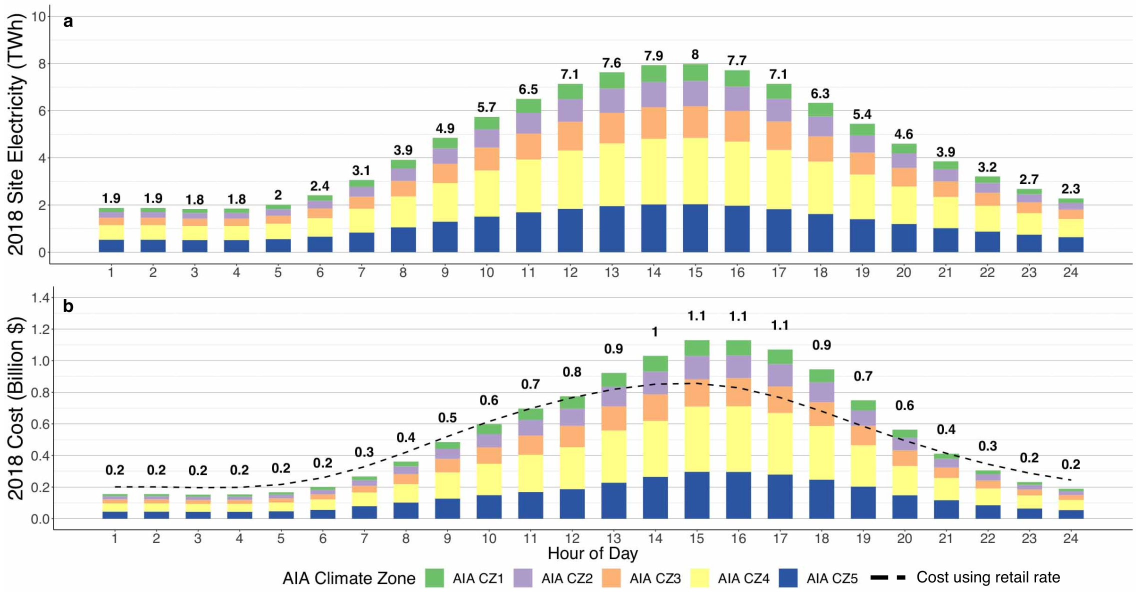

Figure 7.A.5: Hourly summer season (May-September) space cooling electricity and cost totals for commercial buildings by climate zone in 2018. Figure elements and interpretation are the same as in Figure 7.2.

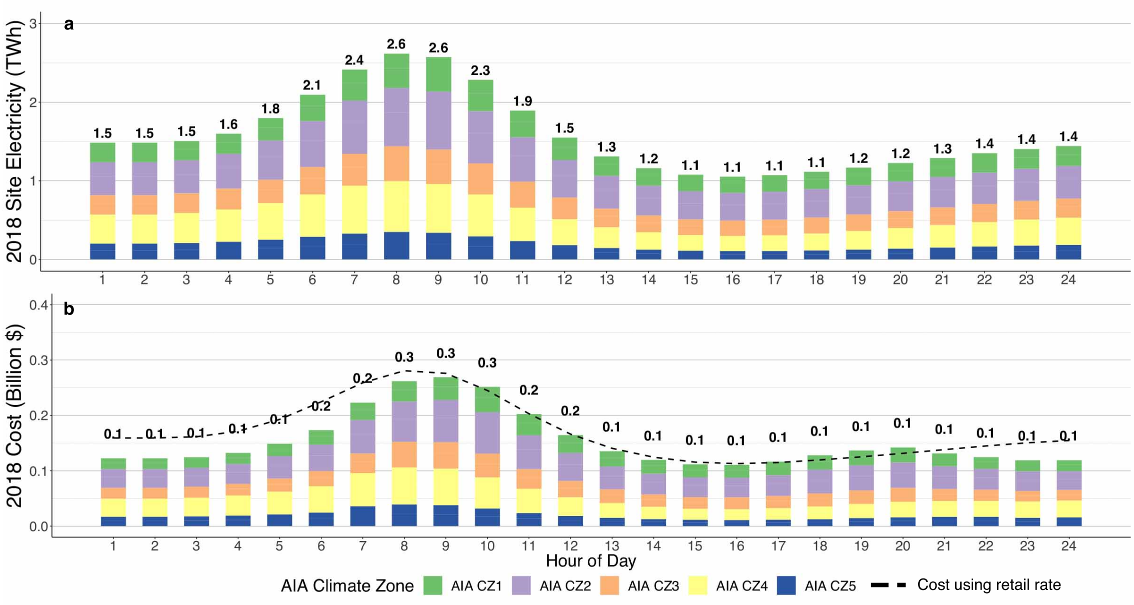

Figure 7.A.6: Hourly winter/intermediate season (October-April) space heating electricity and cost totals for commercial buildings by climate zone in 2018. Figure elements and interpretation are the same as in Figure 7.2.

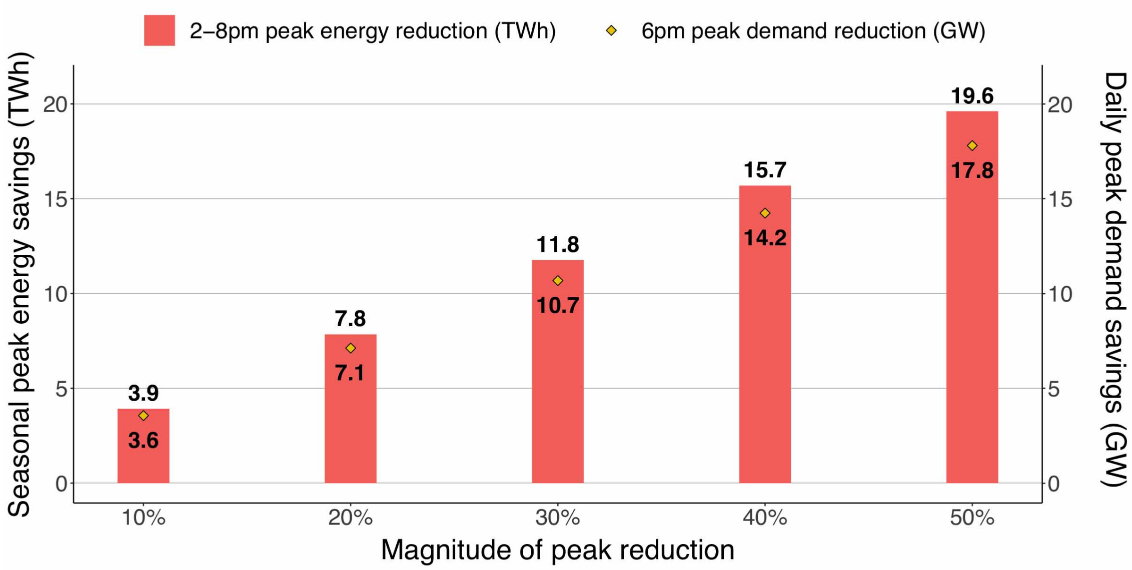

Figure 7.A.7: Summer season peak energy and demand reduction potential for cooling in commercial buildings at different magnitudes of peak energy use and demand reduction. Figure elements and interpretation are the same as in Figure 7.4.

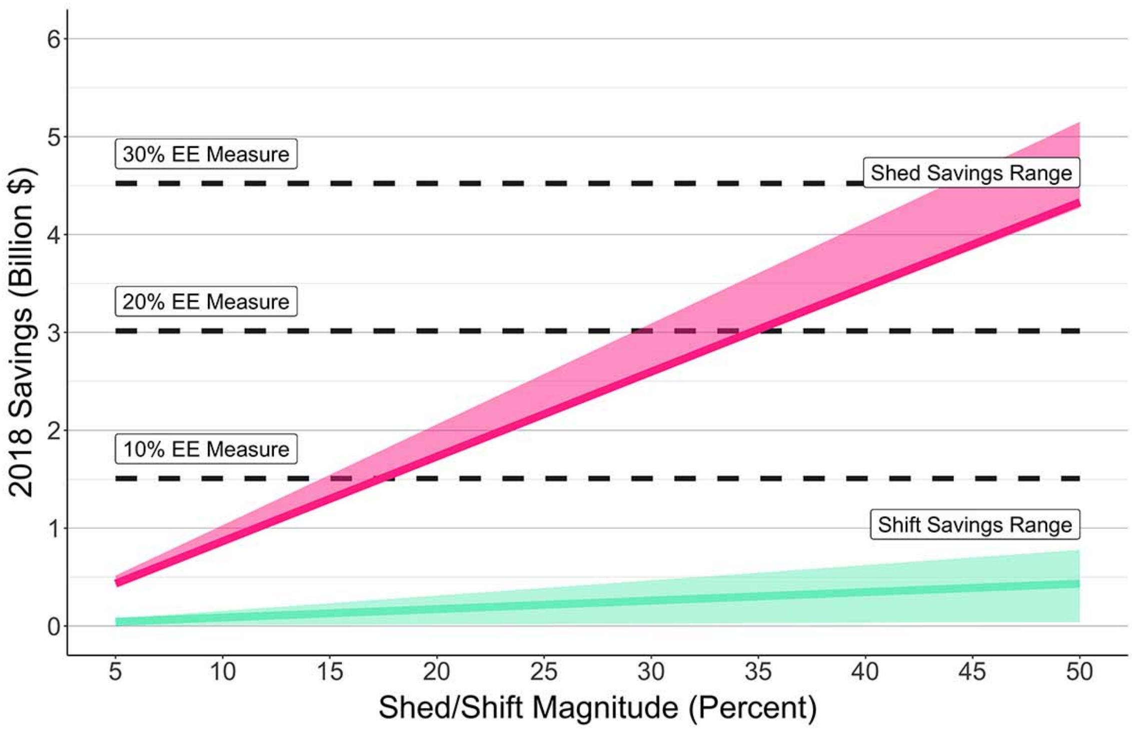

Figure 7.A.8: Technical cost savings potential for static efficiency and dynamic flexibility measures applied to commercial cooling loads. Figure elements and interpretation are the same as in Figure 7.5.

7.B Overview of Scout's analysis approach

Scout analyses center on building energy conservation measures (ECMs) that improve upon the unit-level energy performance and/or operation of a comparable baseline building technology or operational approach. ECM definitions are defined primarily by five attributes: applicable baseline market segment(s), year of market entry/exit, unit-level energy performance, installed cost, and lifetime.

Given one or more ECM definitions, Scout calculates the total impact of each ECM's adoption by consumers or organisations on metric M in year y under adoption case s, $M_{y,s}$:

$$M_{y,s} = \sum_{z}^{Z} \sum_{b}^{B} \sum_{f}^{F_b} \sum_{u}^{U_{b,f}} \sum_{t}^{T_{b,f,u}} \sum_{v}^{V} M_{z,b,f,u,t,v,y,s}\, a_{z,b,f,u,t,v,y,s} \tag{7.10}$$

$$M_{z,b,f,u,t,v,y,s} = M^{\text{base}}_{z,b,f,u,t,v,y} - M^{\text{ecm}}_{z,b,f,u,t,v,y,s} \tag{7.11}$$

where $M^{\text{base}}_{z,b,f,u,t,v,y}$ is the baseline quantity of impact metric M attributable to climate zone z, building type b, fuel type f, end use u, technology t, building vintage v, and projection year y, $M^{\text{ecm}}_{z,b,f,u,t,v,y,s}$ is the same quantity after ECM adoption, Z is the set of all climate zones affected by the ECM, B is the set of building types affected by the ECM, F_b is the set of fuel types for building type b that are affected by the ECM, $U_{b,f}$ is the set of end uses for building type b and end use u that are affected by the ECM, $T_{b,f,u}$ is the set of technologies for building type b, fuel type f, and end use u that are affected by the ECM[10] and $a_{z,b,f,u,t,v,y}$ is a competition adjustment factor that removes overlaps between the applicable baseline market of the ECM and competing ECMs in a portfolio. Impact metrics include primary energy use ($E_{y,s}$, $E^{\text{base}}_{y,s}$, $E^{\text{ecm}}_{y,s}$), CO_2 emissions ($C_{y,s}$, $C^{\text{base}}_{y,s}$, $C^{\text{ecm}}_{y,s}$), and operating costs ($\psi_{y,s}$, $\psi^{\text{base}}_{y,s}$, $\psi^{\text{ecm}}_{y,s}$).

Scout's baseline data are drawn from the U.S. Energy Information Administration (EIA) 2018 Annual Energy Outlook (AEO) Reference Case Scenario (U.S. EIA, 2018*d*). The

[10] For the heating and cooling end uses, an additional distinction between supply-side and demand-side technologies is made, where an air-source heat pump is an example of the former and a highly insulating window is an example of the latter.

data used in this study come from the 2018 version of AEO, which includes projections of U.S. residential and commercial building energy use out to 2050. Baseline energy data in AEO are provided with the granularity needed to focus on specific energy use segments, attributing energy use to particular combinations of climate zone, building type and vintage, fuel type, end use, and technology type. Nevertheless, additional adjustments are necessary to translate the AEO data to the final impact metrics M introduced in Eq. 7.10.

First, total primary energy use is defined as follows:

$$E^{\text{base}}_{z,b,f,u,t,v,y} = E^{\text{ref}}_{z,b,f,u,t,v,y}\, SS_{f,y} \tag{7.12}$$

where $E^{\text{ref}}_{z,b,f,u,t,v,y}$ is the total delivered (or 'site') baseline energy use for a given stock segment (z, b, f, u, t, v) and year y in the AEO Reference Case[11] and $SS_{f,y}$ is the site to source (or 'primary') energy conversion factor for baseline fuel type f and year y. Note that while conversions from site to source energy are necessary for subsequent estimations of emissions and cost baselines, detailed below, all electricity use estimates presented in this paper reflect site energy given the objective of estimating efficiency and flexibility impacts on total electric demand rather than total electric generation.

Site-to-source conversion factors $SS_{f,y}$ for all non-electric fuels are assumed to be unity given the consumption of these fuels on-site. For electricity, the site-to-source factor $SS_{f=\text{elec},y}$ is calculated as:

$$SS_{f=\text{elec},y} = \frac{\left(\Omega^{\text{site}}_y + \Omega^{\text{loss}}_y\right)}{\Omega^{\text{site}}_y} \tag{7.13}$$

where Ω^{site}_y is total delivered electricity for residential and commercial buildings in year y, Ω^{loss}_y is total electricity-related generation, transmission, and distribution losses in year y.[12]

Given a baseline segment's total primary energy use $E^{\text{base}}_{z,b,f,u,t,v,y}$, the CO_2 emissions associated with that energy use, $C^{\text{base}}_{z,b,f,u,t,v,y}$, are calculated:

[11] Drawn from the raw AEO output files 'RESDBOUT.txt' for the residential sector (U.S. EIA, 2018*f*) and 'KSDOUT.txt' for the commercial sector (U.S. EIA, 2018*e*).

[12] Total delivered electricity and electricity-related losses data for the residential and commercial sector are drawn from AEO Summary Table A2 (U.S. EIA, 2018*c*).

$$C^{\text{base}}_{z,b,f,u,t,v,y} = E^{\text{base}}_{z,b,f,u,t,v,y}\, CI_{f,y} \tag{7.14}$$

where $CI_{f,y}$ is the CO_2 emissions intensity for primary energy of baseline fuel type f in year y, calculated as:

$$CI_{f,y} = \frac{C^{\text{ref}}_{f,y}}{E^{\text{base}}_{f,y}} \tag{7.15}$$

where $C^{\text{ref}}_{f,y}$ is the total CO_2 emissions reported for residential and commercial buildings, fuel type f, and year y in the AEO reference case,[13] and $E^{\text{base}}_{f,y}$ is the total primary energy use in residential and commercial buildings for the same fuel type f and year y.

Similarly, the energy costs associated with $E^{\text{base}}_{z,b,f,u,t,v,v,y}$, $\psi^{\text{base}}_{z,b,f,u,t,v,y}$, are calculated as:

$$\psi^{\text{base}}_{z,b,f,u,t,v,y} = E^{\text{base}}_{z,b,f,u,t,v,y}\, FC_{b,f,y} \tag{7.16}$$

where $FC_{b,f,y}$ is the primary energy cost[14] for building type b and baseline fuel f in year y.

The reader is referred to the Supplemental Information section of Langevin et al. (2019) for full details on Scout ECM definitions, baseline data, and calculation methods.

Calculating time-sensitive impacts for energy efficiency and flexibility measures

To calculate the time-sensitive impacts of different energy efficiency and flexibility measures presented in Figures 7.4–7.5 and in Table 7.1, we modify the sub-annual baselines based on percentage reductions to yield new estimates for eqs. 7.6–7.9. For the static efficiency measures (Figure 7.3a), we estimate new electricity load shapes for a 10%, 20%, and 30% load reduction that is applied evenly across all hours. The flexibility measures (Figures 7.3b–7.3c) modify the baselines based on percentage reductions during peak hours or percentage shifts from peak hours to off-peak hours. We estimate new baselines for eqs.

[13]From AEO Summary Table A18 (U.S. EIA, 2018*b*).
[14]From AEO Summary Table A3.

7.6–7.9 for different magnitudes of percentage reductions and for different assumptions regarding the duration of the measure and hourly price intensities ($\Phi^{\text{pr}}_{c,b,s,d,h}$). For the shed and shift measures represented in Figures 7.3b–7.3c, we estimate new $E^{\text{base-st}}_{z,b,u,t,v,y,s,d,h}$ and $\psi^{\text{base}}_{z,b,u,t,v,y,s,d,h}$ for the following scenarios:

- 'Low savings potential': applies the 5th percentile of all price intensities ($\Phi^{\text{pr}}_{c,b,s,d,h}$) and a four-hour measure duration (4–8 p.m.)

- 'Medium savings potential': applies the 50th percentile of all price intensities and a six-hour measure duration (2–8 p.m.)

- 'High savings potential': applies the 95th percentile of all price intensities and an eight-hour measure duration (12–8 p.m.)

For the shift measure represented in Figure 7.3c and Figure 7.5, we estimate new $E^{\text{base-st}}_{z,b,u,t,v,y,s,d,h}$ based on the percentage reductions stipulated above. For this hypothetical measure, we assume no net increase in daily energy load and simply distribute reductions in peak load evenly over the pre-cooling window (which is the previous 4, 6, and 8 hours for the 'low', 'medium', and 'high' scenarios, respectively).

In Table 7.1, we assess specific measures based on reported point estimates of peak reductions from existing studies. For each of these measures, we consider a peak window of 4–8 p.m. For the shift measure, in particular, we translate results from Turner et al. (2014), which uses building simulation tools to investigate peak cooling energy demand reduction potential of mechanical pre-cooling in residential buildings across 15 cities in representative U.S. climate zones. The housing prototype used in the simulations is a lightweight wood-frame home with thermal performance typical of new construction. For more specific assumptions regarding building size, occupancy, and thermal performance, we direct readers to the study.

The results presented in Turner et al. (2014) include total annual peak period cooling energy reductions (Table 7 in the study), fractional peak period energy reductions (Table

8), and peak-to-penalty energy ratios (Table 9) for mechanical pre-cooling across the representative climate zones. We use these published estimates to yield national weighted averages for the following:

- $|\Delta E_p|$, the peak period energy reduction (which can be expressed as a percentage, R),
- ΔE_c, the increase in off-peak cooling energy (energy penalty) above that required to match the on-peak decrease in cooling energy, due to operating equipment under a pre-cooling schedule,
- Γ, the 'peak-to-penalty' energy ratio, which is calculated as: $\Gamma = \frac{|\Delta E_p|}{\Delta E_c}$.

Specifically, we first calculate the energy penalty ΔE_c from Tables 7 and 9 in the study using the listed 'Short' pre-cooling strategy (corresponding to our 4–8 p.m. peak and 12–4 p.m. pre-cooling windows) and a thermostat set point of 23.3°C (from a base set point of 25°C). Next, we use these estimates of ΔE_c and the values for ΔE_p from Table 7 to calculate a national weighted average for Γ of 3.19. Using the same pre-cooling strategy and set-point assumptions, we then use Table 8 to calculate a national weighted average for R of 65% reduction during the peak period using the pre-cooling strategy. We apply these Γ and R values to arrive at new totals for (7.6)–(7.9), where Γ is used to calculate an energy penalty that is distributed evenly across all hours outside the peak window, and R is used to calculate both the percentage reduction during the peak window and the resulting increase in energy during the pre-cooling window. Finally, we use these new measure baselines to yield total energy, cost, and emissions impacts for the shift measure.

7.C Time-sensitive adjustment figures

Figure 7.C.1 plots the median and 5th/95th percentile of weekday and weekend load fraction shapes from the EPRI database in the peak (summer) season, across pre-2004 NERC regions by end use (figure columns) and building type (figure rows). Here, load fractions, $\phi_{r,b,u,d,h}$,

are calculated by normalising the raw peak season hourly demand intensities from the EPRI library by total daily demand intensity for each combination of pre-2004 NERC region r, building type b, end use u, day type d, and hour h as follows:

$$\phi_{r,b,u,d,h} = \frac{L_{r,b,u,d,h}}{\sum_{h=1}^{24} L_{r,b,u,d,h}} \tag{7.17}$$

where L_t is the raw average demand figure for each hour h, under the given region, building, end use, and day type combination.

In many cases, little variation is observed in load shapes across pre-2004 NERC regions after normalisation (indicated by the difference between the 5th/95th percentile lines from the median); in such cases, while the *magnitude* of hourly demand may vary regionally, the *shape* of the demand remains similar. Figure 7.C.1 also shows important differences between residential and commercial load shapes (top and bottom rows, respectively), with residential loads generally pushed later in the day, and between weekday and weekend load shapes, with weekend load shapes often appearing flatter.

Figure 7.C.2 summares daily TOU rate shapes across the U.S. by month and building type, showing the median weekend and weekday rate shape for each month/building type combination as well as the 5th, 25th, 75th, and 95th percentiles of the weekday rate shape. Across building types, median weekday rate shapes peak in July and are at their lowest in January; peak hourly rates occur at roughly 6 p.m. for both residential and commercial customers, though the 95th percentile of the weekday shapes shows a somewhat earlier peak for the commercial sector.

Figure 7.C.3 shows the hourly marginal CO_2 emissions factor curve for each season and NERC region, as well as the median national curve. In general, there is muted variation in marginal emissions factors across the day, with the most variation seen in the summer (nationally, from approximately 500 to 700 kg CO_2 per MWh). In most cases, emissions factors in the overnight and early-morning hours (e.g., off-peak hours) are higher than those estimated for daytime hours—a trend that is particularly noticeable in the summer season—likely reflecting the use of gas-fired generators on the margin to manage the morning

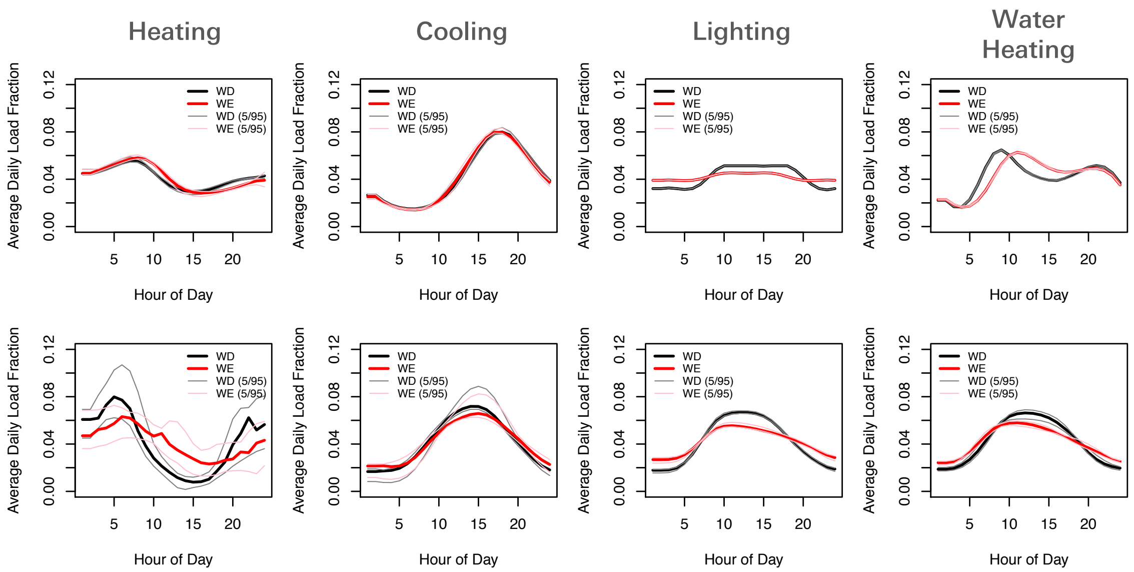

Figure 7.C.1: Median and 5th/95th percentiles of average hourly load fractions from the EPRI End Use Load Shape database (EPRI, 2012) across pre-2004 NERC regions (U.S. DOE, 2017), on weekdays ('WD') and weekends ('WE') in the peak season (May-September), for residential buildings (top row) and commercial buildings (bottom row). Load fractions indicate the portion of total daily load attributable to a given hour for a given building type, day type, and end use combination. Load shapes for other end uses (e.g., refrigeration, appliances) and the off-peak season (October-April) are available from the EPRI library but not shown here.

and evening ramps versus the use of coal during the overnight hours with low demand (Siler-Evans et al., 2012).

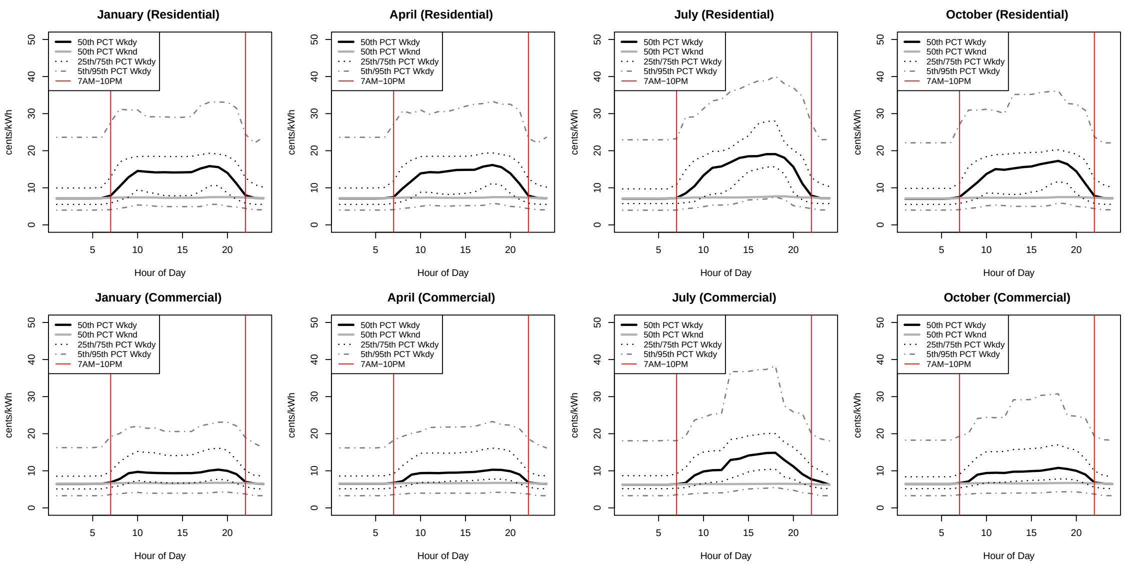

Figure 7.C.2: Median national weekday ('Wkdy') and weekend ('Wknd') hourly electricity prices under time-of-use rate structures in the URDB (NREL, 2018), broken out by month (columns) and building type (rows). Also plotted are the 5th, 25th, 75th, and 95th percentiles of the hourly price shapes for weekdays, as well as a reference 7 a.m.-10 p.m. time window, representing typical hours of building occupancy. Price shapes for this figure are calculated based on an aggregated analysis of data that span all utilities listed with residential or commercial TOU rates in the URDB (287 utilities and 1251 rates for residential buildings, 306 utilities and 2016 rates for commercial buildings).

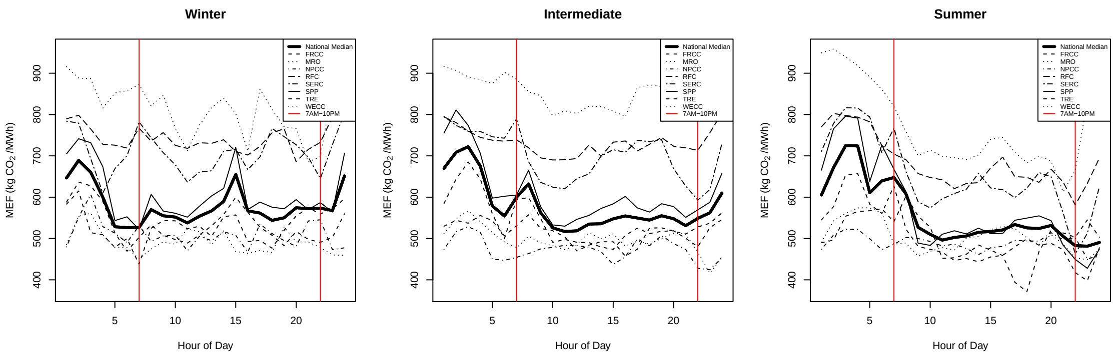

Figure 7.C.3: Hourly marginal CO_2 emissions by post-2011 NERC region (U.S. DOE, 2017) and season as determined by a previous analysis of data from the EPA Continuous Emissions Monitoring System (CEMS) (Siler-Evans et al., 2012). Also plotted are the national median of the seasonal marginal emissions shapes, as well as a reference 7 a.m.-10 p.m. time window, representing typical hours of building occupancy. Seasons are defined as winter (December-February), intermediate (March-April, October-November), and summer (May-September).

Chapter 8

Discussion

8.1 Introduction

This chapter summarises the findings of the thesis' papers and discusses how their results link together to answer the research questions set forth. Section 8.2 reviews key findings from the four papers and discusses methodological, empirical, and policy implications for the thesis as a whole along with reflections on the thesis' contributions. Section 8.3 identifies limitations in addition to those discussed in the individual papers. Section 8.4 outlines avenues for further research.

8.2 Summary and implications of research

This thesis set forth three overarching research questions, and each of these is further informed by the sub-questions outlined in the four papers. The following sections summarise the findings pertinent to these research questions and place the results in the context of what is known from the literature. To aid the reader, Tables 8.1–8.4 present high-level findings specific to each of the papers' research questions.

8.2.1 Key methodological findings

Table 8.1 presents a summary of the methodological findings from the four papers. These findings and their implications for energy research are discussed further in the subsequent paragraphs.

The first paper of this thesis presents an introduction to regularisation methods and their use in modelling annual residential electricity consumption with sparsity (Hastie

	Paper 1	Paper 2	Paper 3	Paper 4
Methodological research question	How can regularisation methods improve variable selection efforts and interpretability in statistical models of residential electricity consumption?	By how much does inclusion of time-use activity data in models of daily and intra-day electricity consumption improve the explanatory power of these models?	What approaches can improve outcomes for clustering high-resolution residential electricity load profiles for the aim of identifying peak-period usage clusters?	How can a time-sensitive framework inform like-for-like comparisons of the savings potential of energy efficiency and flexibility measures in the U.S. building sector?
Summary of methodological research findings	Regularisation methods show lower prediction error and greater sparsity in models of residential electricity consumption, thus demonstrating how they improve upon conventional variable selection techniques (e.g., stepwise regression) in both respects	Inclusion of time-use data in models of intra-day electricity use increases model explanatory power by an average of 0.09 in R^2.	Clustering cumulative rather than raw electric load profiles allows for straightforward use of Euclidean distances in cluster analysis while still capturing the ordered, temporal nature of load profiles.	By disaggregating annual building electricity use to sub-annual temporal resolutions, a time-sensitive framework enables more accurate assessments of the comparative energy, cost, and emissions savings of various energy efficiency and flexibility measures.

Table 8.1: Summary of the four papers' methodological findings.

et al., 2015). It demonstrates how these techniques can yield highly predictive statistical models while also identifying influential variables from large sets of possible predictors, thus improving model interpretability through variable selection (James et al., 2013). When applied to a large sample of California households, two of the sparse regularisation methods, elastic net and group lasso regression, show lower prediction error than the methods that have most often been used for variable selection in these kinds of studies (e.g., stepwise regression). Furthermore, all of the sparse regularisation methods included in the analysis yield simpler models than the stepwise regression model. This finding highlights that these methods are useful for excluding irrelevant or non-significant predictors without sacrificing predictive accuracy.

This result has several methodological implications for energy research: first, given the costs of data collection, prioritising which variables to include in predictive models can reduce the expense of studies of residential electricity consumption. Regularisation methods that can identify influential predictors while excluding irrelevant variables are especially useful in this domain because there exists an almost limitless number of factors that potentially shape residential electricity demand (Hsu, 2015). Second, regularisation methods that perform variable selection can identify those predictors that show strong associations to electricity consumption and thus may be particularly effective to target with educational or technological interventions.

Given that the first paper in this thesis is the only one to compare model results across both regularisation and alternative approaches for variable selection, several other points relevant to this comparison are given here. As was discussed in Chapter 3, the negative attributes of subset selection methods such as stepwise regression make these methods problematic given associated issues with inflated model coefficients and R^2 values, falsely low standard errors and narrow confidence intervals, and an increased likelihood of Type I errors (Harrell, 2001). Despite the fact that these methods perform variable selection and so may be useful for improving model interpretability by identifying significant predictors and excluding irrelevant ones, there are numerous reasons to be sceptical of their results. In the first paper, the type of stepwise regression method used is backward elimination using AIC for model selection, which has several advantages over forward selection (Mantel, 1970). Because this method first fits the full model, it is likely to retain more predictors than forward selection, which may explain the higher number of non-zero coefficients in Table 4.8. Still, the results clearly show that several regularisation methods (elastic net and group lasso) have lower prediction error than the stepwise regression while also achieving more parsimonious models. Beyond the theoretical reasons for avoiding subset selection methods, the methodological results from paper 1 confirm that shrinkage methods can perform better in terms of both prediction and interpretability.

It is also true, however, that the gains in predictive accuracy demonstrated in the first paper are slight, and the first paper concludes that this is likely due in part to the fact that even the full models do not appear to overfit the data (possibly because n is much larger than p), so the performance benefits of regularisation are not as strongly showcased as they could be. But as they are applied to increasingly complex data in the second and third papers, the advantages of these methods in addressing key analysis issues become clear.

The second and third papers analyse data sets that combine conventional dwelling, appliance, and socio-demographic variables with occupant time-use data. These papers aim to understand more occupant-specific factors that shape consumption on a daily basis and during the hours of peak demand. The usefulness of regularisation methods for discovering complex activity-electricity use relationships is clear: in a similar way that they identify key variables in models of annual electricity consumption in the first paper, they do so in models of intra-day electricity consumption in the second paper and peak-period consumption patterns in the third paper. In terms of improving model interpretability via variable selection, the second paper shows that lasso achieves a reduction from 46 to 13 predictors in models of daily average de-minned electricity consumption, and the third paper shows that elastic net is able to reduce a predictor set of 177 hour-activity variables to just 22 to predict peak-period cluster membership. In the third paper, in particular, typical approaches such as OLS would not work in this high-dimensional setting, so regularisation techniques are the most appropriate methods for achieving interpretable yet accurate models in this context.

The key variables identified through the analyses in papers 2 and 3 are useful for predictive purposes, as well. While paper 1 does not use the same train-test framework as does the first and third papers, the identification of key activity variables and their associations with daily and four-hour averages of electricity consumption improves these models by increasing their explanatory power, again showcasing the ability of regularisation methods to select important variables and exclude irrelevant ones. In the third paper, the regularisation method is applied in a new prediction task—that of binary classification—and shows predictive accuracy on the test set that is competitive with a non-linear tree-based method. The third paper also demonstrates a new approach for pre-processing electric load profiles to show how Euclidean distances can be appropriately used to capture the full shape of a load profile when using cluster analysis. This approach has the benefits of conceptual simplicity and ease of implementation, and it enables straightforward clustering of time-series data.

Returning to the thesis' methodological research question on what statistical learning approaches can improve efforts to characterise residential electricity consumption patterns, the discussion above highlights that regularisation methods can improve upon past efforts to model residential electricity demand using conventional regression and subset selection approaches—both in terms of model predictive accuracy and in terms of model interpretability. Beyond these two criteria, regularisation methods also prove useful in addressing several additional analysis issues. To show that these methods have indeed improved upon past efforts, a summary of the key analysis issues at stake and how these are or are not addressed in each of the first three papers is presented in Table 8.2.

This table highlights criteria that can be used to evaluate the performance of the statistical learning methods applied in the thesis and shows how and under what circumstances regularisation methods provide an improvement. As has been discussed at length, regularisation methods deliver improvements related to model interpretability via more parsimonious models and prediction accuracy. Next, regarding the issues of high-dimensional data and overfitting, both of which bear relevance to the aims of model interpretability and prediction, the first and second papers do not as strongly showcase the suitability of regularisation methods to address these, but the third paper certainly does, highlighting how conventional approaches cannot even be used. Similarly, regarding the issue of multicollinearity, this problem is not apparent in the first two papers but is in the third paper, wherein the choice of regularisation method is guided by this analysis issue (Zou and Hastie, 2005).

Though Chapter 3 includes a comparison of the positive and negative attributes of non-linear methods as well as linear methods, these are only compared in paper 3, and the results show that random forests performs better than regularisation in the classification task, supporting the claims made in Table 3.1 that non-linear methods often outperform other approaches in prediction. Yet these methods do not deliver results that are as

Analysis issue	**Summary of how regularisation methods address analysis issues**		
	Paper 1	**Paper 2**	**Paper 3**
Model interpretability (variable selection)	Elastic net, group lasso, and adaptive lasso reduce predictor set from 58 to 11–21 (60–80% reduction), yielding a simpler model than that selected using stepwise regression (see Table 4.8, p. 143)	Lasso reduces predictor set from 46 to 13 (70% reduction) (see Table 5.4, p. 192)	Elastic net reduces predictor set from 177 to 22 (87% reduction) (see Table 6.3, p. 251)
Prediction accuracy	Elastic net and group lasso improve model prediction accuracy over stepwise regression (see Table 4.8, p. 143)	Test set prediction error not explicitly tested, though model containing all variables (no regularisation) shown to have higher prediction error than lasso model (see Figure 5.4, p. 190)	Elastic net shows slightly lower prediction accuracy than random forests; no comparison with conventional regression methods possible because $p \gg n$
High-dimensionality and overfitting	Because $n \gg p$, models do not appear to overfit training data (though AC-only model without regularisation does show some evidence of overfitting—see Figure 4.A.2, p. 166)	Because n is not much larger than p, least squares fit exhibits high variance, so lasso is used for variable selection (see Figure 5.4, p. 190)	Because $p \gg n$, there is no unique least squares coefficient estimate; OLS cannot be used so elastic net and random forests are used instead
Multicollinearity	Multicollinearity effects are not observed	Multicollinearity effects are not observed	Multicollinearity present due to large predictor set with high incidence of near-zero variance predictors; elastic net improves handling of multicollinearity (Zou and Hastie, 2005)

Table 8.2: Summary of how and when the regularisation methods applied in this thesis have improved efforts to characterise residential electricity consumption by addressing key analysis issues.

straightforward to interpret as those yielded by sparsity methods.

In summary, while non-linear methods might perform better than linear methods with regularisation in complex prediction tasks, regularisation techniques show competitive prediction results with far greater gains for model interpretability because they can perform variable selection. Thus, when induction is an objective, regularisation methods improve upon both conventional regression and non-linear methods. Furthermore, when the number of predictors is equal to or greater than the number of observations, as well as when multicollinearity effects are present, regularisation methods and non-linear methods may perform similarly in prediction (and both will likely perform better than conventional or subset selection methods), but regularisation methods will again yield more interpretable models.

Returning to the claims made in the introduction and in paper 1, variable selection is a key issue in energy research because of the richness of the built environment, wherein there exist such a large number of potential factors that influence consumption, and because of the costs of data collection, which leads to small underlying samples and thus a higher likelihood that models will overfit. In these contexts, regularisation methods can improve upon both conventional and non-linear methods across a number of key criteria. This does not mean they are a panacea, however, and their use should be guided by the specific issues and objectives of the analysis. Indeed, these methods do have some limitations and weaknesses that have been discussed in Chapter 3 and will be elaborated on further in Section 8.3.2. But given how rarely they are applied in energy research, and simultaneously how often they are the subject of studies in statistical learning, they should be considered by far more energy researchers in statistical models of energy demand.

This section has so far avoided discussion of the fourth paper and its methodological implications. That is because it does not use the same bottom-up statistical analysis framework as do the first three papers, so its methodological findings and implications are somewhat different. The fourth paper develops a methodology for deriving hourly estimates of national-scale U.S. building end-use electricity consumption and the seasonal and geographic variability of these from annual baselines. This methodological improvement facilitates the comparison of electricity saving measures with time-sensitive impacts alongside conventional energy efficiency measures. Using a time-sensitive valuation framework to estimate the cost and emissions saving potential of various efficiency and flexibility measures is critical for positioning the buildings sector as a key resource for demand-side mitigation, and this paper details an approach other energy researchers and policy analysts can use to better assess the benefits of demand-side measures in the buildings sector.

	Paper 1	Paper 2	Paper 3	Paper 4
Empirical research question	Which structural and occupant factors are the strongest predictors of annual electricity consumption in California households?	How do the relationships between structural and occupant factors, including occupants' daily activities, and electricity consumption vary as a function of the time of day in U.K. households?	What are the activity-based drivers of peak-period electricity consumption patterns in U.K. households?	How do efficiency and flexibility measures compare in terms of their time-sensitive energy, cost, and emissions savings potential in the U.S. building sector?
Summary of empirical research findings	Structural factors related to dwelling size, type, and efficiency are key predictors, but so are occupant-related variables, such as energy curtailment behaviours.	Appliance ownership and occupant socio-demographic factors show the strongest relationships to daily electricity use patterns, but these relationships vary based on time of day; housework, food, and recreational activities are key for understanding patterns of electricity use throughout the day	Food-related activities, as well as laundry practices, TV/computer use, and socialising, are key predictors of evening peak-period electricity consumption patterns.	Efficiency measures continue to show high cost and emission savings potential under a time-sensitive framework, whereas the potential of flexibility measures is dependent on assumed electricity rate structures and measure magnitude and duration.

Table 8.3: Summary of the four papers' empirical findings.

8.2.2 Key empirical findings

Table 8.3 presents a summary of the empirical findings from the four papers, which are informed by the thesis' empirical research question on what factors explain electricity consumption in residential buildings at various temporal resolutions.

As found in the first paper, which investigates drivers of annual electricity consumption in California households, key variables include structural factors, such as the size, type, and efficiency of the dwelling, ownership of specific devices and appliances, primarily EVs and AC, and occupant-related socio-demographic and behavioural factors, such as household occupancy, income, age, and self-reported energy saving behaviours. Importantly, for this study sample, pro-environmental attitudes and rates of energy literacy are much less important for explaining variations in annual electricity consumption.

These findings confirm much previous research on the primary determinants of annual electricity consumption, especially regarding the influence of the size of home and number of occupants (Brounen et al., 2012; Huebner et al., 2016; Jones et al., 2015; Kavousian et al., 2013; Wyatt, 2013), the type of dwelling (Huebner et al., 2016; Wyatt, 2013; Yohanis et al., 2008), and household income (Cramer et al., 1985; Brounen et al., 2012; Yohanis et al., 2008). The finding that occupant attitudes and literacy are less important in explaining consumption patterns is also supported by previous research (Bartiaux and Gram-Hanssen, 2005; Brounen et al., 2013; Huebner et al., 2016; Vringer et al., 2007). While they have been included less frequently in similar studies, occupant energy curtailment behaviours have shown associations with lower annual electricity consumption in previous research (Wallis et al., 2016).

These results confirm that policies or programs to improve the thermal efficiency of residential buildings are critical for achieving further reductions in usage (Brown and Li, 2019; Méjean et al., 2019). They also suggest, however, that interventions focused on helping occupants develop habitual energy saving behaviours can reduce electricity consumption, especially for cooling-related behaviours (Sonnenschein et al., 2019). In comparison, interventions focused on raising awareness of environmental consequences from energy consumption or instilling pro-conservation attitudes and energy literacy may not achieve similar reductions. It is true that these elements may be antecedents to energy conservation behaviour, and while it cannot be completely ruled out that attitudes and literacy are an underlying driver of energy saving behaviour, the finding that these are not strongly correlated in the study sample and the consistent selection of behavioural rather than attitudinal variables in the regularisation methods tested suggests that behaviours rather than attitudes show a stronger association to reductions in electricity use. The oft-cited fact that people consume energy for the services it provides rather than for the energy itself is instructive here (Janda, 2011); this detachment among occupants suggests that no matter how well-intentioned one might be when it comes to saving energy, it is the actual practice of using less that counts (Shove, 2010; Shove and Walker, 2014).

The results from the second paper show that daily average electricity consumption in U.K. households is shaped primarily by dwelling, appliance ownership, and socio-demographic factors. Size and type of home, EV ownership, and number of occupants are found to be influential (just as they were for explaining annual consumption). Type of electricity tariff, ownership of other large appliances such as gas boilers, power showers, and TVs/computers, and cat ownership are also selected as significant predictors of daily average consumption. Again, many of the findings from the second paper confirm previous research on drivers of daily variations in electricity demand. Size and type of home, number of occupants, and appliance ownership have all been studied extensively and found to associate with increased daily demand (Anderson et al., 2016; Firth et al., 2008; Kavousian et al., 2013; Khan et al., 2019; McLoughlin et al., 2012).

In models of four-hour average de-minned electricity consumption, where baseload demand is removed from each household, several activity variables are found to be influential in explaining these variations. Activities related to caring for the home, eating and meal preparation, and recreation are factors that explain variations in consumption across four-hour periods throughout the day.

These empirical results from the second paper are further confirmed in paper 3, which takes a different analytical approach to predicting and understanding consumption patterns in U.K. households by using cluster analysis and classification methods. The paper finds that food-related activities are important predictors of consumption patterns throughout the evening hours of 5 to 9 p.m. Additional activities that explain consumption patterns during these hours are laundry-related activities, especially washing machine use, as well as TV/computer use, socialising, and the timing of arriving home in the evening. Strategies to reduce demand during peak hours can take these activities into account to achieve greater impact; for instance, interventions can be tailored to households based on activities that may constrain or complement demand shifting.

Though occupant time use and activities have been studied less frequently, previous research supports this thesis' findings on activities that are associated with increased energy usage, such as food and meal preparation (De Lauretis et al., 2017; Druckman et al., 2012; Jalas and Juntunen, 2015), laundry and other housework activities (Anderson, 2016; Palmer et al., 2013; Widén et al., 2009), and sleeping or resting (Druckman et al., 2012; Jalas and Juntunen, 2015). Anderson and Torriti (2018) find that changes in the timing of peak demand in the U.K. have had little to do with changes in media use, which contradicts the finding from paper 3 that TV/computer use is a strong predictor of peak-period consumption.

The empirical findings from the first three papers on what factors shape demand at different temporal resolutions deliver insights for strategies for demand reduction and response. Occupant activities are of special interest in these analyses, and the findings highlight the types of activities that might be particularly useful to consider when designing interventions to shift demand. Of course, these findings are constrained to specific geographies given the areas in which the studies are undertaken. Still, they identify key factors that are related to patterns of electricity consumption, many of which have not been properly examined in previous research. The following section will return to this point to further discuss the policy implications of these findings.

Though it does so via an alternative approach, the fourth paper also presents evidence on the the factors that shape electricity consumption in U.S. residential buildings. It finds that thermal end uses (space cooling and heating) are responsible for the largest share of electricity consumption in residential buildings and contribute substantially to peak demand, accounting for over 45% of usage during the peak hour. The paper further shows that electricity use varies regionally and seasonally. Summer electricity consumption for space cooling nationally in the U.S. is five times larger during the peak hours of 5 to 6 p.m. than during early morning hours (6 to 7 a.m.) and is primarily driven by cooling in southern climates. These findings show which major electric end uses drive consumption patterns at an annual as well as hourly temporal resolution in the U.S. building sector.

While these findings are not directly comparable with those of the first three papers because of the different modelling approaches used and the different factors considered, the results certainly support the first three papers' findings that thermal and electrical appliances are substantial drivers of residential electricity consumption at annual and hourly temporal resolutions.

8.2.3 Key policy implications

Table 8.4 presents the policy research questions and summarises the policy implications of each of the four papers. These are framed by the thesis' policy research question on how a more detailed understanding of the temporal aspects of residential electricity consumption can inform strategies and policies for energy efficiency and flexibility interventions.

The first paper presented in this thesis has key policy implications for achieving further reductions in electricity consumption in California households. Interventions to reduce demand should be informed by empirical evidence on which factors are substantial contributors to demand. In the case of Palo Alto, the prevalence of large, detached homes—many of which have AC and EVs—appears to be a substantial contributor to electricity consumption patterns among the city's residents. Given these empirical findings, energy efficiency interventions should focus on improving the thermal performance of the existing residential building stock through encouraging energy audits, retrofit campaigns, or uptake of efficient appliances and devices, as these were found to be underutilised by the study sample. Municipal policies could encourage denser urban planning and smaller home sizes to reduce usage, but these might be politically challenging to implement. This highlights the fact that identifying a key driver of consumption does not necessarily mean it is a useful lever for change. Still, knowledge of the factors that influence consumption can inform which subsets of the housing stock or the population to target with energy efficiency interventions. In Palo Alto, AC and EV ownership is associated with increased consumption, so targeting these households with smart charging technology or smart thermostats might prove useful. Beyond technology-focused interventions, behavioural campaigns and trials to encourage

	Paper 1	**Paper 2**	**Paper 3**	**Paper 4**
Policy research question	How does identifying key drivers of annual residential electricity consumption inform energy efficiency policy design in California?	How can understanding which factors influence intra-day variations in household electricity consumption inform policies for demand flexibility in the U.K.?	How does evidence on the relationships between occupant activities and electricity consumption during peak hours provide insights for targeted interventions to reduce U.K. peak demand?	How does quantifying the time-sensitive impacts of U.S. building electricity use inform investments in national energy efficiency and flexibility technologies?
Summary of policy implications	Policy efforts should focus on improving the thermal performance of homes and increasing uptake of energy efficiency programmes, developing energy saving or smart charging interventions targeted to households with AC and EVs, and encouraging energy curtailment behaviours among occupants.	Flexibility policies should be based on the knowledge that the effect of dwelling and household variables on electricity consumption varies as a function of time of day and that activity-related factors that influence electricity consumption vary from day to day	The design of pricing, incentive, and behavioural DSR interventions should take into account the timing of activities related to food preparation and meal times, laundry practices, and TV/computer use given the strong associations of these to peak-period electricity demand patterns	Investments in building technologies should focus on thermal end uses, especially space cooling, given their substantial contribution to annual and peak building electricity demand; high-efficiency HVAC systems with controls that enable peak shedding and load shifting yield significant energy, cost, and emissions savings.

Table 8.4: Summary of the four papers' policy implications.

energy saving behaviours among residents (especially the elderly or those who live in large houses with AC) might be more favourable and less costly, even though they may not yield as large of savings. These kinds of policy comparisons must be guided by evidence on the factors that influence consumption.

While efficiency policies and interventions are clearly important for reducing energy use and emissions, policies that aim to increase flexibility in the residential sector, whether through technology deployment, economic and pricing programmes, or behavioural interventions, must consider how electricity consumption varies throughout the day. The second and third papers in this thesis provide empirical evidence on the factors that shape daily and peak-period patterns of consumption, which can inform DSR programme design. For instance, given the associations found between washing machine use and increased consumption in the early evening (around 5 p.m.), a utility could use this information to encourage or support adoption of smart controls for washing machines that can shift their use to off-peak hours. If these are deployed alongside TOU rates that charge lower rates for off-peak consumption, both the utility and its customers can benefit from energy and cost savings by shifting demand. If deploying this technology is not feasible for cost or other reasons, a less expensive behavioural intervention could be designed to deliver text messages or similar notifications to occupants that using the washing machine during early evening hours is especially costly for them and for the electricity system. This is just one example, and although it may seem straightforward and obvious to many, people often do not understand or are unaware of the benefits of reducing their consumption during times of peak demand, so helping them do so remains an important aim (Darby and McKenna, 2012).

More generally, a detailed understanding of what energy *is used for* in homes can inform socio-technical approaches that do not conceive of reducing energy demand simply as a technological problem but instead understand it as a challenge requiring deep insight into the preferences and practices of individuals, the technologies they use, and the social context for their electricity consumption (Shove and Walker, 2014; Geels et al., 2017). Such an approach might be able to envision radical, systemic change in the energy system and then further investigate how public policy can deliver this change. Understanding the social aspects of consumption are essential to this objective, so providing evidence on the activities or practices that result in energy use can inform not only specific, targeted policies but also perhaps much more far-reaching, systemic ones, too.

While this is true, it is still the case in much of the world that policies are guided by assessments of their economic benefits or the economic value of other benefits, such as reductions in carbon emissions. For this reason, it is important to quantify the potential benefits of energy efficiency and flexibility technologies to guide policy and investment.

The fourth paper's findings do just that. First, they provide evidence that thermal end uses are substantial contributors to both annual and peak electricity demand, meaning developing measures to reduce or shift demand from these end uses can achieve significant cost and emissions savings. Second, by assessing the time-sensitive impacts of several efficiency and flexibility measures applied to residential space cooling, the paper presents an illustrative application of how to compare these measures on a like-for-like basis, finding that existing cooling efficiency measures could yield an estimated 57.3 TWh of seasonal electricity savings, 29.8 GW of summer daily peak demand savings, \$8.1 billion USD in cost savings, and 26.4 $MtCO_2$ of avoided emissions. Flexible cooling measures that reduce consumption during peak hours can save 8.4 TWh, 14.4 GW, \$1.5 billion USD, and 3.8 $MtCO_2$. Instead of reducing this demand, shifting it to off-peak hours via a pre-cooling technology could greatly increase daily peak demand savings to 62.4 GW, though such a measure would entail a slight increase in seasonal electricity use of around 6% of total summer season residential cooling demand. A measure combining the best efficiency and flexibility features would accrue the largest seasonal electricity (63.1 TWh), daily peak demand (39.8), cost (\$9.2 billion USD), and emissions (29.1 $MtCO_2$) savings. This finding is intuitive and demonstrates the additional benefits that can be achieved when coupling efficiency measures with flexibility controls.

These findings can inform federal investments in building technology R&D by clearly illustrating the cost and emissions savings from efficiency and flexibility technology adoption, thus making a strong business case for why these types of measures should be prioritised in decarbonisation pathways. Further, the findings have useful insights for thinking further about how enabling policies such as time-varying tariffs could improve the cost effectiveness of deploying these technologies. Given that the savings of the flexibility measures, in particular, are highly dependent on assumed electricity rate structures, proper design of these could facilitate more investment in and adoption of demand-side flexibility resources.

The question of cost effectiveness is an important one that is not assessed in the fourth paper, primarily due to the absence of reliable cost estimates for new flexibility technologies.

But future work to quantify which efficiency or flexibility measures are 'worth it', either assessed across the entire building stock or for individual consumers, can further guide R&D policy and investment decisions. Given that the analysis in paper 4 is mostly focused on the benefits to consumers (who are the ones who would see the cost savings from reducing or shifting demand when on a TOU rate), the question of cost effectiveness has related questions about who will provide the upfront investments for these new technologies and, if it is the individual consumer, what sort of payback period and return on investment can be expected. A quick calculation based on the results of paper 4 suggests that if the cost savings from the static cooling efficiency measure were assessed on a per household basis, the seasonal savings would amount to around $68 per household, assuming around 120 million U.S. households (U.S. Census Bureau, 2019). Would this justify the cost of installing a new, high-efficiency cooling system, even for purely rational homeowners? Would this incentive be meaningful in the context of an average U.S. residential electric bill that totals $1,400 per year (U.S. EIA, 2017)? These are important questions that require further consideration. While it is true that a move to time-varying prices in the residential sector can improve the cost-benefit calculations for peak reductions, it is also the case, as has been discussed previously, that electricity demand is not highly price responsive, and there is no guarantee that savings will be realised for consumers on time-varying electricity tariffs. Novel approaches in DSR programme design are thus needed to ensure that the benefits of building technology adoption do, in fact, reach the end users.

Finally, the fourth paper also sheds light on the nuance of designing emissions reduction policies given specific geographical or energy system contexts. The figure showing hourly variations in end-use residential electricity, cost, and emissions (Figure 7.1, p. 281) indicates that the period when electricity supply is most carbon intensive in the U.S. is during nighttime when demand is low. This is perhaps counterintuitive to some, but as paper 4 explains, it is most likely due to coal being on the margin during nighttime and early-morning hours, whereas gas-fired generators are supplying peak demand (Siler-Evans et al., 2012). This finding highlights that policies and interventions to reduce energy use might not be aligned with those that are designed to reduce carbon emissions (e.g., Graff Zivin et al., 2014). Effective policy design must take into account regional or local energy system dynamics given that the carbon intensity of electricity generation can vary greatly across different geographical contexts as well as within the same geography across time due to the evolving nature of energy supply (Cubi et al., 2015; Khan et al., 2018).

8.2.4 Summary of research outputs

These four papers collectively develop and apply new methods to better predict electricity consumption patterns at annual, daily, and hourly temporal resolutions. They also identify the key factors driving these patterns and in doing so provide new evidence on important levers for delivering more efficient and flexible electricity demand in the residential sector.

Because particular attention in the papers is given to occupant-specific factors, especially occupant activities, one important output of this work is a publicly-available data set of high-resolution electricity and time-use data, which researchers can use to improve or validate simulated models of electricity demand or conduct more detailed investigations of time-use and energy-use relationships. This is the first data set of its kind in terms of its parallel records of high-resolution electricity and occupant time-use data.

A second output of the thesis research is a national baseline for time-sensitive efficiency and flexibility valuation in the U.S. building sector, which is also being made available to energy researchers. This baseline enables others to conduct their own analyses of how different efficiency and flexibility measures compare in terms of their potential economic and emissions benefits.

The four papers included in this thesis thus provide new analytical methods, new empirical evidence, new data sources, and a new valuation framework, all of which aim to inform policies and strategies for rapid decarbonisation of building electricity use.

8.2.5 Reflections on research contributions and novelty

As one of the three research questions for this thesis focuses on identifying factors that explain patterns of electricity consumption, the informed reader may question the contributions this work makes to the already expansive literature on determinants of residential electricity demand. Many of the factors highlighted as key drivers of consumption, especially those related to occupant socio-demographic and physical dwelling characteristics, have been studied extensively in the literature across different geographical contexts (Jones et al., 2015). For that reason, several further reflections on the novelty of this work and its key contributions to the literature are provided here.

First, while it is true that there is a large number of published papers that take a similar empirical approach of statistically modelling residential electricity consumption, often using linear regression and its variants, to quantify the positive or negative effects of different types of household and dwelling factors, it is also true that the majority of these studies focus on annual or monthly electricity consumption. Primarily due to smart meter data availability, studies investigating daily or hourly variations in consumption and the factors influencing these are far less common. These types of studies are becoming increasingly common with access to these new data, however, but there still remains an evidence gap on how different types of household, dwelling, and occupant factors influence intra-day or hourly patterns in demand. Access to interval electricity data from households presents new opportunities to learn about what influences electricity consumption at higher temporal resolutions, and this knowledge becomes increasingly important for policy goals related to peak demand reduction and demand flexibility.

Second, in part due to this proliferation of and access to high-resolution data, key analysis issues related to high-dimensional data and overfitting are becoming increasingly common in energy research. In this context, it is important to reiterate the need for methods that predict well out of sample but also yield more interpretable results by excluding irrelevant variables. While methods that primarily focus on building energy prediction and forecasting are useful for a variety of purposes, so is learning where policy efforts should be focused to deliver further reductions in electricity demand. For the latter aim, variable selection is critical. Most of the approaches in building energy analysis are dominated either by conventional regression and model selection approaches (OLS, subset selection) or by complex, non-linear statistical learning methods (e.g., artificial neural networks, support vector machines). The interpretation of results in both cases can be challenging, albeit for different reasons that have been explained thoroughly. Finding a middle ground between these two approaches—one that maintains strong predictive accuracy while also improving model interpretability—is essential for deepening insights into complex energy usage dynamics in the residential sector while discovering new levers to change consumption patterns. This thesis has clearly demonstrated how a class of methods that are rarely applied in energy research can improve upon existing approaches to characterise patterns of demand and identify the influential factors shaping them. It has also provided a nuanced examination of key analysis issues, such as high-dimensional data, overfitting, and multicollinearity, and has used this discussion to highlight the circumstances under which statistical learning with sparsity is most likely to improve these efforts. In this way, the thesis makes novel contributions to the fields of energy analysis and modelling.

Third, regarding the thesis' empirical contributions, although the first paper's findings may seem somewhat obvious at first glance, it should be reiterated that the paper investigates numerous factors that are much less frequently studied, such as occupant energy curtailment behaviour and energy literacy. Still, the paper's main contributions stem from the introduction of the regularisation methods discussed above, which become increasingly applicable to the high-resolution energy and activity data sets that are the subject of papers 2 and 3. Here is where the empirical findings are more central. The empirical evidence base on the relationship between occupant activities and electricity demand is substantially less developed than the evidence linking dwelling and socio-demographic factors to consumption (Grunewald and Diakonova, 2018; Powells et al., 2014). The data and evidence provided in papers 2 and 3 provide new understanding of both how traditionally-studied factors vary in their influence on consumption throughout the day as well as how new activity factors affect consumption. These outputs stress the importance of considering what energy *is used for* in homes (Gram-Hanssen, 2014; Shove and Walker, 2014), and this contribution is certainly novel in the context of statistical analyses of electricity demand.

Fourth, in terms of the contributions to effective policy design, this thesis has provided numerous insights that can inform the development of targeted interventions to reduce or shift demand in the residential sector. Beyond these, the thesis has also provided novel comparisons of how energy efficiency and flexibility measures compare in terms of their energy, cost, and emissions savings. In terms of informing low-carbon policy and strategy, the results from the fourth paper are perhaps the most useful. National-scale building technology policy and investments must be guided by quantitative assessments of the kind presented in paper 4, and as the paper makes clear, previous analyses that account for time-sensitive variations in building energy use and its associated costs and emissions are significantly lacking. Filling this gap and sharing a new analysis framework that can facilitate further assessments of the benefits of energy efficiency and flexibility interventions is another novel contribution of this thesis.

8.3 Research limitations

While this thesis makes novel contributions to building energy research via new analytical methods, data sources, and evidence on factors influencing consumption, there are numerous limitations across these areas. This section discusses limitations to the data and analytical methods used. Some of these are mentioned in the individual papers, but the focus here is discussing broader limitations that are pertinent to the thesis as a whole.

8.3.1 Data limitations

Sample representativeness and size

Though this limitation is discussed thoroughly in papers 1–3, its recurrence in this thesis means it is worth repeating here: the study samples on which the analyses and results of the first three papers are based are not representative and include several sources of bias as a result of the convenience sampling method used. An abbreviated list of these include an over-representation of higher income households, larger dwellings, more educated occupants, more conservation-oriented occupants, and somewhat more energy literate occupants (though comparing energy literacy for the samples to the population of interest is difficult given a lack of large-scale studies investigating energy literacy). Given findings from past research, it is almost certain that these biases influence the results to some extent. Quantifying this 'influence' is difficult, but the papers aim to ensure transparency around this issue by making a comparison between the study sample and the population of interest to clearly show where these biases exist.

Several other types of bias are present: the issue of social desirability response bias is present in studies where survey instruments are the primary data collection method. In the surveys developed for this research, there are opportunities for respondents to give socially desirable rather than factually accurate responses to survey questions. A similar bias may be prevalent in the activity data collected for papers 2 and 3, as activities are all self-reported. Some checks for accuracy and consistency are discussed in the papers, but there remain limitations to how much the factual accuracy of activity reporting can be assessed.

Sample size is arguably a lesser issue than sample representativeness when it comes to developing inferential models (unless the sample size is too small to achieve the necessary statistical power for a given significance level), but given that this thesis research does not delve as much into statistical inference (see Section 8.3.2), the larger issue related to sample size has to do with the analytical methods that are used. The modelling techniques demonstrated throughout this thesis are, in the literature, often applied to larger data sets than those collected here. This means that some issues of model stability and uncertainty are present. We return to a discussion of this issue in Section 8.3.2.

Finally, regarding the specific study samples, not much previous discussion has been given to the geographical differences among these. The choice of study area for the individual papers was guided mainly by opportunities for data access or partnership with industry or policy groups. The thesis' methodological research question is one that could reasonably be applied to data from many different geographies. Indeed, I conclude in several of the papers that the methodological approaches are applicable in other contexts, as well. The thesis' empirical research question, however, is necessarily limited by its geographical scope, especially given how different the electricity systems are in the U.S. and U.K. (as well as compared to other countries). While there is much opportunity for comparative analyses of electricity use patterns and their determinants in different regions, this was not one of the objectives of this research. For this reason, and also given the aforementioned sample representativeness limitations, the papers contextualise study results within the geographical regions in which they are undertaken.

Factors not considered

While the papers in this thesis have investigated a wide range of factors that potentially influence residential electricity demand, they by no means capture all the possible factors that shape consumption patterns. Due to restrictions to the number of survey questions that could be asked, and given the specific research aims of this thesis, factors had to be prioritised in advance of data collection.

Broadly, two related aims directed the design and implementation of the surveys: first, survey instrument design was guided by past studies in order to include factors that have consistently shown strong relationships to residential electricity consumption so that these could be controlled for in the models. Second, the research design was guided by a specific focus on occupants and their role in delivering demand reduction and response. This is why, for instance, behavioural curtailment factors were prioritised over detailed appliance data in paper 1. Similarly, the proposed method for time-use data collection in papers 2 and 3 is tailored such that more energy-relevant activities are studied. Still, there are several cases where limitations to the factors considered in the models might reduce their predictive accuracy while also limiting their ability to highlight influential variables.

This may be particularly relevant in the third paper, where only activity variables were included in the predictive models (rather than socio-demographic data, which was readily available from the survey). The decision to exclude socio-demographic and dwelling variables in the third paper was intentional—there are numerous previous studies that have examined the relationships between load profile shape and socio-demographic variables (McLoughlin et al., 2015; Rhodes et al., 2014; Viegas et al., 2016), whereas few if any have investigated the predictive power of activities alone. This was one of the objectives of paper 3, which is why the analysis is constrained to include only activity data. It should be mentioned, however, that the conference paper that preceded this journal publication did examine the relationships between socio-demographic, dwelling, and appliance factors and load profile shape, though it did not do so using the same statistical methods as in paper 3 (Satre-Meloy et al., 2019). Future work could apply these methods to the survey data to see how well they can predict evening peak-period usage profiles.

A specific area pertinent to residential electricity demand that may appear absent from this thesis is data on appliances and behaviours related to thermal end uses. Indeed, as paper 4 shows, from a national building energy modelling perspective, these are primary drivers of consumption. In the first three papers, these are covered in a rather indirect fashion. In paper 1, the survey asked about AC ownership but not source used for heating, whereas in paper 2 the survey asked about ownership of either a gas boiler or heat pump. These questions could have been simplified by asking respondents what source of energy they use for heating, as it would be expected that electric heating would have a substantial impact on electricity consumption, especially in the U.K. These variables are included in the models presented in paper 2, but interpreting their effects is not straightforward.

Furthermore, occupants are not asked to share detailed heating and cooling behaviours (for instance, through the activity recording devices). Specific data on heating and cooling behaviours could provide a better understanding of how these fit into daily routines and subsequently influence electricity use. One way to address this issue would be to focus specifically on non-heating electricity use by excluding households with electric heating or cooking appliances (e.g., Huebner et al., 2016). Overall, in the samples studied in this thesis, the incidence of electric heating is generally low, so excluding these households would not likely change results significantly.

A related limitation is that the electricity consumption data included in the first three papers are not weather normalised. In the first paper, the study sample includes households all located in the same city, so weather-related effects are assumed to be even across the study sample. In the second and third papers, however, the samples consist of households spread throughout the U.K. In this context, especially given that households participate in the study in different seasons, weather normalisation would normally be appropriate. This step is not taken, however, because most of the study sample in papers 2 and 3 use gas for space heating and do not have AC, so weather-related effects on electricity consumption are again assumed to be minimal. Still, this is a limitation that could be addressed in future analyses, especially those that aim to better understand heating and cooling-related activities.

Decarbonising electricity is a critically important challenge for the low-carbon transition, which is why this thesis has focused specifically on characterising electricity demand. But other fuels deserve attention, too, especially the use of gas for cooking and heating. While the current design of the electricity recorder device is not suitable for collecting gas interval data, this is a limitation that will be addressed in future research.

Time-sensitive assessment framework limitations

A final data-related limitation is specific to paper 4: the data sources compiled to enable time-sensitive estimates of energy, cost, and emissions in U.S. buildings each has limitations.

Paper 4 comments briefly on these and discusses how they will be addressed in future research, but a few other points are mentioned here. The load shape data used to disaggregate national-level electricity consumption estimates to hourly totals come from a database maintained by the Electric Power Research Institute (EPRI, 2015). These data are from simulations conducted in the early 2000s that are based on numerous assumptions. The credibility of these data were questioned during peer review. While the paper acknowledges the shortcomings of this data source, given it is based not on empirical data but on simulations for different building prototypes, it is arguably the best available data source for hourly end-use load shapes that are specified regionally and seasonally, which is central to the analysis conducted in paper 4.

One unexpected finding from the work to search for and integrate the data sources for paper 4 is how little empirical data actually exists for hourly end-use variations in electricity consumption. As mentioned in the paper, the U.S. DOE has recently funded a 3-year project scheduled to conclude in 2021 to develop and validate end-use load profiles for the entire U.S. building stock for all 8,760 hours of a typical meteorological year (U.S. DOE, 2013). These data will be incorporated into the Scout analysis framework in the future, but this is a limitation at present to the data used to disaggregate electricity end-use consumption to a sub-annual basis. This absence of credible empirical data on drivers of hourly electricity consumption patterns is, in a way, the motivation for much of the work undertaken in this thesis.

8.3.2 Methods limitations

Statistical inference

While the statistical methods applied in this thesis have advantages for energy research, they have limitations as well. The primary limitation is the inability to generalise results to a wider population due to the absence of much statistical significance testing in the papers. It should be mentioned that given the non-random nature of the samples used, statistical inference would not generally be appropriate in this context, anyway. This is one of the reasons the research papers are less concerned with confidence intervals and significance testing and more focused on identifying key variables in predictive models of electricity consumption.

The other reason is one that is mentioned in the papers but will be reiterated here given its relevance: most inferential constructs (such as confidence intervals and p values) do not generally exist for regularisation methods because of the way these algorithms search for strong relationships in the data in order to select key variables; under these circumstances, assessing the uncertainty of these estimates and declaring them statistically significant or not is challenging. As mentioned in paper 1, there is an active literature on post-selection inference using the lasso and other regularisation techniques (Lockhart et al., 2014; Taylor and Tibshirani, 2015; Lee et al., 2016). There is no question that inference is important for assessing the uncertainty of model coefficients and the significance of results in studies of residential electricity consumption. Though the literature on statistical learning with sparsity is beginning to develop techniques for inference (see references above), the studies included in this thesis do not apply these, and this is a limitation of its methodology. Even given the inability to generalise beyond the boundaries of the data sets analysed, however, the papers' findings have insights that can be used to inform policy, as detailed in Section 8.2.3.

A related limitation is the issue of endogeneity in statistical models. An endogeneity problem occurs when there is a correlation in a model between one of the predictors X and the model's true error term ϵ. What this signals is that there is another variable not included in the model that is correlated with both X and the dependent variable Y. The consequence of this is that the estimated effect between X and Y will be biased and inconsistent. Endogeneity arises due to omitted variables but also from other sources that are not easily measured, such as sample selection and simultaneity (reverse causation). Issues of endogeneity may be particularly relevant to the data and results from the first paper, where the magnitude of the relationship modelled between energy curtailment behaviours and electricity consumption may be due to indirect effects.

There are several approaches to handling endogeneity in regression models, such as the introduction of an instrument variable, which is correlated with X but uncorrelated with ϵ, or using an experimental research design with random treatment assignment (Bascle, 2008). This research takes neither approach for the same reasons given above—namely, the studies are not concerned with statistical inference but rather the demonstration of sparsity techniques for improving model prediction and interpretability. While acknowledging these limitations, the results are still useful for informing future efforts to understand patterns of electricity consumption using study designs that prioritise causal inference.

Inconsistencies in application of methods

The attentive reader will note that paper 2 does present confidence intervals and significance levels in models of daily and intra-day electricity consumption, even after using lasso regression for variable selection. To be transparent, these were added upon recommendation of one of the paper's peer reviewers. While the better approach would be to use the post-selection inference methods discussed above, because the models in paper 2 presented OLS coefficients rather than lasso coefficients, the post-selection inference method described in Lockhart et al. (2014) was not used to compute confidence intervals and significance levels. A similar approach is taken in Huebner et al. (2016).

The other key difference in paper 2 is the absence of the train/test framework used in papers 1 and 3. While cross-validation is used to tune the penalty parameter for the regularised regressions (just as it is in papers 1 and 3), in paper 2 the focus is less on model predictive accuracy and more on the comparison of model coefficients for different factors across four-hour periods of the day. It would be equally appropriate, however, to take a similar train/test approach for each of the four-hour models and compare them based on RMSE or a similar criterion rather than R-squared. The train/test framework is the more common (and arguably more appropriate) approach within statistical learning disciplines (though not necessarily within energy demand research).

Some additional inconsistencies in the way the methods are applied are apparent in the analytical differences between papers 1–2 and paper 3. While papers 1 and 2 model a continuous dependent variable (annual/daily average consumption), paper 3 models a binary dependent variable (early- or late-peaking cluster). This is primarily a matter of preference and interest in studying and applying a wide range of statistical learning methods, including those in both regression and classification settings. That is not to say, however, that a cluster-classification approach would not be appropriate for characterising annual electricity consumption nor that a regression approach is not appropriate for characterising peak-period consumption. Both types of approaches are useful and have been demonstrated effectively in the literature. The motivation for incorporating classification algorithms in paper 3 is to show how the regularisation methods developed in papers 1–2 can be used in this setting, as well.

Several other inconsistencies in terms of the application of statistical learning in this thesis include the following: 1) the advanced lasso methods (group and adaptive lasso) are not applied in papers 2 and 3 because these papers are less concerned with comparing regularisation methods and their extensions and more focused on picking the method that is most appropriate given the data and modelling task; 2) while multiple imputation is used for missing data in paper 1, this method is not used in papers 2 or 3 given the lower incidence of missing data, the expectation that its impact on results would be negligible, and the need to keep the analysis from becoming overly complex; 3) statistical interactions between variables in the regression and classification models presented are not considered, primarily because these become increasingly challenging to interpret (especially for greater than two-way interactions); previous work, however, has shown the effectiveness of regularisation methods for handling interactions while improving predictive power (Hsu, 2015), so interactions should be considered in future research.

The experimental nature of paper 3

Returning to paper 3, it is important to note the experimental nature of both the modelling approach and data preparation in the context of statistical models of electricity consumption. To the authors' knowledge, this is the first paper to use occupant activity data in predictive models of peak-hour load profile clusters. The use of regularisation methods in this context is also novel and experimental. As the paper makes clear, the results are dependent on the myriad pre-processing and evaluation steps taken throughout both the cluster and classification analyses. The motivation, in part, is to show the usefulness of the methods that are developed in the first two papers of the thesis for handling complex, high-dimensional data. But combining these multifaceted analysis steps to assess the relationships between occupant activities and peak-hour electricity consumption is an untested approach.

Furthermore, substantial research and analyses were undertaken in preparation of this paper that were not included in the final submitted version. This research and analysis mostly focused on reviewing past methodological approaches for load profile clustering, comparing these empirically using our data, and assessing the stability and robustness of clustering results under different approaches and assumptions. While the paper was originally written to include these results, it was decided this additional work should be more carefully presented in a separate paper that focuses specifically on methodological considerations. This paper is currently being written for submission to a methods-focused energy journal.

8.4 Future work

8.4.1 Statistical learning and characterising electricity demand

This thesis presents detailed applications of statistical learning methods for characterising electricity demand in the hope that these methods will continue to be used for energy demand research. In particular, it highlights the advantages of statistical learning with sparsity for modelling complex human-energy dynamics. This is an actively developing field of statistics research (Hastie et al., 2015). Future work should consider how new developments in statistical learning with sparsity may be beneficial for energy demand research, especially for handling increasingly large and complex energy and activity data sets that face challenges related to interpretability and overfitting.

As interpretability in models of electricity demand is key for delivering insights into various strategies for demand reduction and flexibility, an exciting area of active research is interpretability in non-linear methods, especially deep learning. Closing the gap between data prediction and knowledge extraction for these methods is an important avenue of research if these highly complex methods are to be applied for practical research purposes (Vellido et al., 2012). There is a groundswell of interest in deep learning interpretability, and this work will hopefully improve efforts to achieve further gains in our understanding of the factors influencing electricity use.

This thesis describes several different methods for modelling building electricity use, including bottom-up statistical and engineering approaches. While the specific modelling techniques common to these approaches are usually distinct, there is scope for combining the two in hybrid approaches that use elements of both approaches for better prediction and interpretation of building electricity use (Swan et al., 2008). One could envision, for example, an approach that uses detailed building physics simulations to estimate specific end uses, such as space heating or cooling, and then combines these with a statistical learning approach that aims to generalise these relationships to much larger samples. Here, sparsity techniques would also prove beneficial by identifying outputs from the engineering models that are most important for explaining variations in electricity use at the scale of a regional or national housing stock.

A related area where advanced statistical learning methods are being applied in a similar fashion is the increasingly popular approach known as energy demand disaggregation, which is also called non-intrusive load monitoring (Hosseini et al., 2017). Disaggregation refers to statistical approaches to take whole-building power signals and disaggregate them into appliance-specific data. The benefit of this approach is the ability to provide very detailed appliance monitoring and feedback in individual buildings without having to monitor individual plug loads (Carrie Armel et al., 2013). This thesis research has not attempted to disaggregate power signals for households, primarily because the focus has been on occupant activity patterns (rather than just appliance power signatures) and electricity consumption. But there is certainly potential in demand disaggregation for both enhancing and further informing analyses linking time- and energy-use relationships. The temporal resolution of the data collected in papers 2 and 3 is appropriate for disaggregation, and this is a possible avenue for future research. Such work may be particularly informed by the accompanying activity data, as these could be used to validate the results of disaggregation algorithms (Najafi et al., 2018).

Disaggregation approaches are also those that have received the most attention from the private sector given their potential for commercialisation. Already, there are several companies that apply high-frequency load monitoring in households to provide unique insights for energy saving, shifting, safety, and other related services.[1] Providing detailed feedback to occupants for improving household energy management is an exciting avenue for commercial enterprises that may have benefits for efforts to decarbonise electricity demand.

8.4.2 Scaling up time-use data collection

Because one of the data limitations is the size and nature of the samples collected in papers 2 and 3, future research should address this limitation through collecting time-use and electricity use data from a larger, more representative sample of households. An important question that has been asked by several reviewers during peer review for this research is whether it is even possible to scale up the collection of time-use data in parallel with interval electricity data using the device-based approach presented in papers 2 and 3.

Scaling up data collection has indeed been harder than anticipated. Several years ago,

[1] Verv is a London-based company that uses a similar approach as the METER project for non-intrusive load monitoring: `https://verv.energy/news/`.

the expectation was that the devised approach would be able to achieve a sample size of 1,000 households in several years of data collection. But the METER team has fallen far short of this goal, mostly owing to the time-consuming nature of data collection (e.g., sending out parcels with electricity and activity-recording devices as well as receiving these back and processing them individually). By the time of this writing, however, the team has made several important changes to the data collection procedures that we believe can enable future scaling up of our sample size.

First, we have vastly simplified the data processing stage such that instead of connecting individual returned devices and downloading the data before uploading to our database, we have created a procedure wherein this all happens automatically via the cloud when the devices are connected to Wi-Fi. This change will speed up the data processing/uploading phase, which has been a bottleneck for data collection.

Second, we have created a mobile application available in the Apple App and Google Play stores so that any future participant with a smart phone can download the activity recorder app and start providing time-use data. This change will both speed up the sending out and returning/processing phases while also doubling the number of devices that we can use for recording electricity data.[2] Participants who want to take part but do not have a smart phone can still participate in the previous way (with a dedicated activity recorder).

This new app-based approach will facilitate longitudinal data collection while also opening up the study data to researchers who would like to use the time-use data collection approach we have developed to study phenomena outside the field of electricity consumption (e.g., the links between health, well-being, and time-use). Furthermore, it is possible that with enough training data of electricity and time-use, we may reach a point where we are able to predict electricity consumption based on activity data with reasonable accuracy. This approach is explored to some extent in paper 3, which uses activity data to predict whether a household is likely to have an earlier or later peak. This possibility is likely

[2] The devices used for recording activities and electricity data are the same but have a different app uploaded for either activity or electricity recording. If we no longer have to send out activity recording devices, we can use all our hardware for electricity recording.

some ways off, however, given the need for much larger data sets to develop more accurate models. But in the future, it might be feasible that activity data alone can give reasonable predictions for household electricity use.

Third, regarding the electricity recording devices, we are exploring ways to extend the battery length for these devices (which is the primary barrier to collecting data for longer than 28 hours). We have piloted the use of external batteries to extend the study duration for an additional 28 hours and are considering other ways to lengthen the study period. We have also worked with researchers in Germany to develop an electricity recording device that is functional with German electricity meters (which do not have exposed mains electricity lines). We are actively exploring other avenues for international research, which would similarly open up opportunities for cultural comparisons of time- and electricity-use relationships.

Finally, the ideal way to scale up data collection for the METER project would be to link our research to existing smart meter data projects in the U.K., specifically the Smart Energy Research Lab (SERL), which aims to provide smart meter data to the energy research community (UCL, 2017). If and when this link is made (for which conversations with SERL are ongoing), we would no longer need to send electricity recorders to participating households. We would simply ask participants for their consent for SERL to share their interval meter data with us, and then we would ask them to use the activity app to record activities for any duration we choose. This combined approach would access much larger (and representative) samples, which would allow the METER project to explore many new avenues for research in the future.

8.4.3 Intervention studies to change demand

While this thesis has principally focused on *understanding* demand, an equally important part of the demand-side research agenda is testing the best approaches for *changing* demand. As the review in Chapter 2 makes clear, approaches from behavioural science and economics, as well as socio-technical studies of energy demand, are key for delivering insights into how to intervene to change consumption patterns. The data generated from papers 2 and 3, especially, can inform future intervention studies.

The METER project has already begun to conduct some of this research. Several pilot studies have been completed to assess different types of interventions for shifting demand. While this type of research is common in energy economics and related disciplines, what the METER project can provide is insight into what actually happens in households when they are asked to reduce or shift demand. This knowledge can help researchers understand the unintended or unexpected consequences of household interventions while also shedding light on how sustainable shifts or reductions might be or how different types of electricity policies might affect households (e.g., dynamic pricing and/or non-monetary incentives). Collecting time-use and electricity data in parallel opens up many future research avenues that did not exist before. These avenues are increasingly important for informing policies and strategies for rapid decarbonisation of residential buildings.

8.5 Chapter summary

This chapter has summarised the key methodological and empirical findings of the four papers and has discussed their policy implications. It has argued that the methods applied in this thesis can improve efforts to model residential electricity consumption while also improving our understanding of the factors that influence electricity consumption patterns in households. It then discussed how this detailed insight can inform the design of policies and strategies for demand reduction and flexibility. It has also summarised the research outputs and commented on how this thesis will contribute to further lines of methodological and empirical inquiry in this field. Finally, it has discussed some of the limitations to this research and avenues for addressing these in future work.

8.6 References

Anderson, B. (2016), 'Laundry, energy and time: Insights from 20 years of time-use diary data in the United Kingdom', *Energy Research & Social Science* **22**, 125–136.

Anderson, B., Lin, S., Newing, A., Bahaj, A. and James, P. (2016), 'Electricity consumption and household characteristics: Implications for census-taking in a smart metered future', *Computers, Environment and Urban Systems* .

Anderson, B. and Torriti, J. (2018), 'Explaining shifts in U.K. electricity demand using time use data from 1974 to 2014', *Energy Policy* **123**, 544–557.

Bartiaux, F. and Gram-Hanssen, K. (2005), Socio-political factors influencing household electricity consumption: A comparison between Denmark and Belgium, *in* 'Proceedings of the 2005 ECEEE Summer Study on Energy Efficiency in Buildings', European Council for an Energy Efficient Economy, Presqu'île de Giens, France, pp. 1313–1325.

Bascle, G. (2008), 'Controlling for endogeneity with instrumental variables in strategic management research', *Strategic Organization* **6**(3), 285–327.

Brounen, D., Kok, N. and Quigley, J. M. (2012), 'Residential energy use and conservation: Economics and demographics', *European Economic Review* **56**(5), 931–945.

Brounen, D., Kok, N. and Quigley, J. M. (2013), 'Energy literacy, awareness, and conservation behavior of residential households', *Energy Economics* **38**, 42–50.

Brown, M. A. and Li, Y. (2019), 'Carbon pricing and energy efficiency: Pathways to deep decarbonization of the US electric sector', *Energy Efficiency* **12**(2), 463–481.

Carrie Armel, K., Gupta, A., Shrimali, G. and Albert, A. (2013), 'Is disaggregation the holy grail of energy efficiency? The case of electricity', *Energy Policy* **52**, 213–234.

Cramer, J. C., Miller, N., Craig, P., Hackett, B. M., Dietz, T. M., Vine, E. L., Levine, M. D. and Kowalczyk, D. J. (1985), 'Social and engineering determinants and their equity implications in residential electricity use', *Energy* **10**(12), 1283–1291.

Cubi, E., Doluweera, G. and Bergerson, J. (2015), 'Incorporation of electricity GHG emissions intensity variability into building environmental assessment', *Applied Energy* **159**, 62–69.

Darby, S. J. and McKenna, E. (2012), 'Social implications of residential demand response in cool temperate climates', *Energy Policy* **49**, 759–769.

De Lauretis, S., Ghersi, F. and Cayla, J.-M. (2017), 'Energy consumption and activity patterns: An analysis extended to total time and energy use for French households', *Applied Energy* **206**, 634–648.

Druckman, A., Buck, I., Hayward, B. and Jackson, T. (2012), 'Time, gender and carbon: A study of the carbon implications of British adults' use of time', *Ecological Economics* **84**, 153–163.

EPRI (2015), Load shape library: Version 3.0, Technical report, Electric Power Research Institute, Palo Alto, CA.

Firth, S., Lomas, K., Wright, A. and Wall, R. (2008), 'Identifying trends in the use of domestic appliances from household electricity consumption measurements', *Energy and Buildings* **40**(5), 926–936.

Geels, F. W., Sovacool, B. K., Schwanen, T. and Sorrell, S. (2017), 'Sociotechnical transitions for deep decarbonisation', *Science* **357**(6357), 1242–1244.

Graff Zivin, J. S., Kotchen, M. J. and Mansur, E. T. (2014), 'Spatial and temporal heterogeneity of marginal emissions: Implications for electric cars and other electricity-shifting policies', *Journal of Economic Behavior & Organization* **107**, 248–268.

Gram-Hanssen, K. (2014), 'New needs for better understanding of household's energy consumption - Behaviour, lifestyle or practices?', *Architectural Engineering and Design Management* **10**(1-2), 91–107.

Grunewald, P. and Diakonova, M. (2018), 'The electricity footprint of household activities–implications for demand models', *Energy and Buildings* **174**, 635–641.

Harrell, F. (2001), *Regression Modeling Strategies: With Applications to Linear Models, Logistic Regression, and Survival Analysis*, Springer Series in Statistics, Springer-Verlag, New York, NY.

Hastie, T., Tibshirani, R. and Wainwright, M. (2015), *Statistical Learning with Sparsity: The Lasso and Generalizations*, Chapman & Hall/CRC.

Hosseini, S. S., Agbossou, K., Kelouwani, S. and Cardenas, A. (2017), 'Non-intrusive load monitoring through home energy management systems: A comprehensive review', *Renewable and Sustainable Energy Reviews* **79**, 1266–1274.

Hsu, D. (2015), 'Identifying key variables and interactions in statistical models of building energy consumption using regularization', *Energy* **83**, 144–155.

Huebner, G., Shipworth, D., Hamilton, I., Chalabi, Z. and Oreszczyn, T. (2016), 'Understanding electricity consumption: A comparative contribution of building factors, socio-demographics, appliances, behaviours and attitudes', *Applied Energy* **177**, 692–702.

Jalas, M. and Juntunen, J. K. (2015), 'Energy intensive lifestyles: Time use, the activity patterns of consumers, and related energy demands in Finland', *Ecological Economics* **113**, 51–59.

James, G., Witten, D., Hastie, T. and Tibshirani, R. (2013), *An Introduction to Statistical Learning*, Vol. 103 of *Springer Texts in Statistics*, Springer New York, New York, NY.

Janda, K. B. (2011), 'Buildings don't use energy: People do', *Architectural Science Review* **54**(1), 15–22.

Jones, R. V., Fuertes, A. and Lomas, K. J. (2015), 'The socio-economic, dwelling and appliance related factors affecting electricity consumption in domestic buildings', *Renewable and Sustainable Energy Reviews* **43**, 901–917.

Kavousian, A., Rajagopal, R. and Fischer, M. (2013), 'Determinants of residential electricity consumption: Using smart meter data to examine the effect of climate, building characteristics, appliance stock, and occupants' behavior', *Energy* **55**, 184–194.

Khan, I., Jack, M. W. and Stephenson, J. (2018), 'Analysis of greenhouse gas emissions in electricity systems using time-varying carbon intensity', *Journal of Cleaner Production* **184**, 1091–1101.

Khan, I., Jack, M. W. and Stephenson, J. (2019), 'Identifying residential daily electricity-use profiles through time-segmented regression analysis', *Energy and Buildings* **194**, 232–246.

Lee, J. D., Sun, D. L., Sun, Y. and Taylor, J. E. (2016), 'Exact post-selection inference, with application to the lasso', *The Annals of Statistics* **44**(3), 907–927.

Lockhart, R., Taylor, J., Tibshirani, R. J. and Tibshirani, R. (2014), 'A significance test for the lasso', *Annals of statistics* **42**(2), 413.

Mantel, N. (1970), 'Why stepdown procedures in variable selection', *Technometrics* **12**(3), 621–625.

McLoughlin, F., Duffy, A. and Conlon, M. (2012), 'Characterising domestic electricity consumption patterns by dwelling and occupant socio-economic variables: An Irish case study', *Energy and Buildings* **48**, 240–248.

McLoughlin, F., Duffy, A. and Conlon, M. (2015), 'A clustering approach to domestic electricity load profile characterisation using smart metering data', *Applied Energy* **141**, 190–199.

Méjean, A., Guivarch, C., Lefèvre, J. and Hamdi-Cherif, M. (2019), 'The transition in energy demand sectors to limit global warming to 1.5°C', *Energy Efficiency* **12**(2), 441–462.

Najafi, B., Moaveninejad, S. and Rinaldi, F. (2018), Chapter 17 - Data analytics for energy disaggregation: Methods and applications, *in* R. Arghandeh and Y. Zhou, eds, 'Big Data Application in Power Systems', Elsevier, pp. 377–408.

Palmer, J., Terry, N. and Kane, T. (2013), Further analysis of the Household Electricity Survey - Early findings: Demand side management, Technical Report 475/09/2012, Cambridge Architectural Research Limited, elementenergy, and Loughborough University, London, United Kingdom.

Powells, G., Bulkeley, H., Bell, S. and Judson, E. (2014), 'Peak electricity demand and the flexibility of everyday life', *Geoforum* **55**, 43–52.

Rhodes, J. D., Cole, W. J., Upshaw, C. R., Edgar, T. F. and Webber, M. E. (2014), 'Clustering analysis of residential electricity demand profiles', *Applied Energy* **135**, 461–471.

Satre-Meloy, A., Diakonova, M. and Grünewald, P. (2019), What makes you peak? Cluster analysis of household activities and electricity demand, *in* 'Proceedings of the 2019 ECEEE Summer Study on Energy Efficiency in Buildings', ECEEE, Presqu'île de Giens, France, pp. 745–755.

Shove, E. (2010), 'Beyond the ABC: Climate change policy and theories of social change', *Environ Plan A* **42**(6), 1273–1285.

Shove, E. and Walker, G. (2014), 'What Is energy for? Social practice and energy demand', *Theory, Culture & Society* **31**(5), 41–58.

Siler-Evans, K., Azevedo, I. L. and Morgan, M. G. (2012), 'Marginal emissions factors for the U.S. electricity system', *Environmental Science & Technology* **46**(9), 4742–4748.

Sonnenschein, J., Van Buskirk, R., Richter, J. L. and Dalhammar, C. (2019), 'Minimum energy performance standards for the 1.5°C target: An effective complement to carbon pricing', *Energy Efficiency* **12**(2), 387–402.

Swan, L., Ugursal, V. I. and Beasuoleil-Morrison, I. (2008), A new hybrid end-use energy and emissions model of the Canadian housing stock, *in* 'International Conference on Building Energy and Environment', Dalian, China.

Taylor, J. and Tibshirani, R. J. (2015), 'Statistical learning and selective inference', *Proceedings of the National Academy of Sciences* **112**(25), 7629–7634.

UCL (2017), 'Smart Energy Research Lab (SERL)', `https://www.ucl.ac.uk/bartlett/energy/research/project-directory/smart-energy-research-lab-serl`.

U.S. Census Bureau (2019), 'U.S. Census Bureau QuickFacts: United States', `https://www.census.gov/quickfacts/fact/table/US/HSD410218#HSD410218`.

U.S. DOE (2013), 'Commercial and residential hourly load profiles for all TMY3 locations in the United States', `https://openei.org/doe-opendata/dataset/commercial-and-residential-hourly-load-profiles-for-all-tmy3-locations-in-the-united-states`.

U.S. EIA (2017), 'Table 5.a - Residential average monthly bill by Census Division, and state', `https://www.eia.gov/electricity/sales_revenue_price/`.

Vellido, A., Martın-Guerrero, J. D. and Lisboa, P. J. G. (2012), Making machine learning models interpretable, *in* 'ESANN 2012 Proceedings', European Symposium on Artificial Neural Networks, Bruges, Belgium.

Viegas, J. L., Vieira, S. M., Melício, R., Mendes, V. and Sousa, J. M. (2016), 'Classification of new electricity customers based on surveys and smart metering data', *Energy* **107**, 804–817.

Vringer, K., Aalbers, T. and Blok, K. (2007), 'Household energy requirement and value patterns', *Energy Policy* **35**(1), 553–566.

Wallis, H., Nachreiner, M. and Matthies, E. (2016), 'Adolescents and electricity consumption: Investigating sociodemographic, economic, and behavioural influences on electricity consumption in households', *Energy Policy* **94**, 224–234.

Widén, J., Lundh, M., Vassileva, I., Dahlquist, E., Ellegård, K. and Wäckelgård, E. (2009), 'Constructing load profiles for household electricity and hot water from time-use data—Modelling approach and validation', *Energy and Buildings* **41**(7), 753–768.

Wyatt, P. (2013), 'A dwelling-level investigation into the physical and socio-economic drivers of domestic energy consumption in England', *Energy Policy* **60**, 540–549.

Yohanis, Y. G., Mondol, J. D., Wright, A. and Norton, B. (2008), 'Real-life energy use in the U.K.: How occupancy and dwelling characteristics affect domestic electricity use', *Energy and Buildings* **40**(6), 1053–1059.

Zou, H. and Hastie, T. (2005), 'Regularization and variable selection via the elastic net', *Journal of the Royal Statistical Society. Series B (Statistical Methodology)* **67**(2), 301–320.

Chapter 9

Conclusions

Accelerating the transition to a low-carbon society to meet ambitious climate change mitigation targets requires rapid decarbonisation of the electricity system. This in turn relies on deep emissions reductions from the residential building sector, a substantial contributor to global electricity demand. This challenge is unprecedented in its scale and will require transformative solutions across the energy system, but there exists a significant gap in evidence on the potential of demand-side solutions for climate change mitigation. Informing demand-side solutions requires first and foremost detailed characterisations of demand, especially of the key structural and social factors that shape consumption patterns.

To contribute to this challenge, this thesis answers three primary research questions. The first is, "What statistical learning approaches can improve efforts to characterise residential electricity consumption patterns?" This question is methodological in scope and is motivated by several challenges facing the building energy modelling community. These include the heterogeneity of building energy demand, which gives rise to an almost limitless number of factors that possibly influence consumption; the cost of collecting detailed observational data on these factors, which means models are often based on small data sets and thus do not predict well with new data; and the interpretability of models, which suffers both when the number of predictors in the model is large and when non-linear methods are used.

This thesis makes an original methodological contribution to the energy modelling literature by showing how approaches from the field of statistical learning with sparsity can address these challenges and improve efforts to characterise patterns of residential electricity consumption. The first three papers introduce and develop regularisation methods for statistical models in both regression and classification contexts. These are applied to increasingly complex data sets that include more than 70 unique socio-demographic, dwelling, appliance, and occupant attitude/behaviour variables as well as more than 60 time-use activity variables. Modelling such a large number of predictive variables as well as such a variety in variable type is not feasible using conventional statistical methods. Regularisation techniques are especially helpful in this context. The papers show how the application of methods that deliver sparse solutions can reduce the number of predictors by 60–85%, thus yielding much sparser, more interpretable models that do not overfit the data and maintain predictive accuracy that is generally consistent with results demonstrated in the literature.

While the application of regularisation in statistical models is the main methodological contribution of this thesis research, the papers also demonstrate the use of several other statistical learning approaches to characterise electricity consumption patterns. The third paper, in particular, shows how several methodological improvements to cluster analysis in the unsupervised learning context can improve efforts to segment households by the full shape of their electric load profile, which is important for understanding temporal variations in electricity consumption.

Turning to the empirical focus of the thesis, the second primary research question is, "What factors explain electricity consumption in residential buildings at annual, daily, and hourly temporal resolutions?" This empirical research question is motivated by the need for further evidence on which structural and occupant-related factors should be prioritised when developing interventions to reduce or shift demand. The former outcome of demand reduction is necessary to limit the extent of the supply-side transformation to renewables and deliver further emissions reductions while this transition is underway. Demand reduction during peak times is especially important for limiting both power sector emissions and the costs of electricity supply. The latter goal of responsive demand is increasingly important in energy systems where variable renewable power contributes a large percentage of electricity generation and thus where demand-side flexibility in residential buildings is essential for limiting the costs and easing the technical challenges of integrating renewables.

This thesis research contributes to the evidence base supporting both of these key objectives for electricity system decarbonisation. Moving beyond analyses that focus strictly on aspects of the physical dwelling, which have been the dominant focus of energy models in past research, this thesis places building occupants and their role in shaping electricity demand at the centre of its analyses. To do so, it collects and analyses time-use data, which have been used for electricity demand modelling in previous research but have never before been collected in such a way as to permit analyses of the direct relationship between time use and electricity use in households. In other words, this thesis research enables much more focused investigations of what electricity *is used for* in homes. Understanding this simple question is fundamental to both demand reduction and flexibility in the residential sector.

The empirical findings from the papers first highlight that a mixture of structural and socio-demographic factors are key to explaining annual electricity use in U.S. and U.K. households, accounting for around 30–40% of variance explained in statistical models. These factors include dwelling size and type, presence of insulation, ownership of specific electronic devices and appliances, and socio-demographic factors indicating number of occupants, household income, and age. These factors all show positive associations to electricity consumption, confirming findings from previous research. Several factors that have not been studied as extensively, such as occupant curtailment behaviours, show negative associations to annual electricity consumption.

Daily average electricity consumption is shaped by many similar factors, especially dwelling size and type, EV ownership, and number of occupants. Ownership of large appliances and TVs/computers also explains increased daily usage. The influence of these factors on electricity consumption is not consistent throughout periods of the day, however. Morning average consumption is shaped by different factors than evening consumption, and this finding highlights that consumption varies as a function of time of day. Occupant time-use activities are essential for understanding time-of-day-dependent electricity consumption.

Activities related to housework, eating and meal preparation, and recreation are associated with increased usage, but their relationship to average electricity consumption similarly varies across different times of day. Housework appears especially relevant for electricity demand in the evenings, whereas recreation activities have stronger associations with consumption in the mid-morning and afternoon.

Food-related activities and meal preparation, especially the use of energy-intensive cooking appliances, show strong associations to patterns of electricity consumption during evening peak hours (5 to 9 p.m.). Other key activities that explain patterns of evening peak-period consumption include washing machine use and laundry activities, socialising, and TV/computer or other personal device use. The finding that these activities show promise for predicting whether or not a household has an evening consumption profile that peaks earlier (around 5 p.m.) or later (around 7 p.m.) suggests that these must be better understood in the context of peak demand reduction strategies.

Regarding the policy focus of the thesis, the third primary research question is, "How does a more detailed understanding of the temporal aspects of residential electricity consumption inform energy efficiency and flexibility interventions?" This research question aims to explore how deeper insight into the time-varying aspects of electricity consumption and the key factors underpinning these can be used to inform more effective efficiency and flexibility interventions.

The policy implications of the papers' findings can be summarised as follows: first, given that dwelling and appliance factors explain a substantial share of variations in residential electricity consumption on both long and short timescales, policy efforts must include improving the thermal performance of buildings and the efficiency of appliances and consumer devices. Second, based on empirical evidence that suggests occupant activities are similarly important for explaining patterns of electricity consumption throughout the day and especially during times of peak demand, the design of efficiency and flexibility policies that rely on price-, incentive-, or behaviour-based interventions must take this evidence into account. Doing so is likely to improve the effectiveness of these policies by enabling more targeted interventions; a deeper understanding of what might be an appropriate incentive or nudge for certain households as opposed to others will likely deliver higher savings at lower costs while also mitigating some of the negative distributional consequences, low magnitudes of response, and high participant attrition rates of blanket approaches to reducing or shifting demand.

Beyond informing the design of more effective interventions, this thesis research also has implications related to national-scale building technology policy. These implications come from the thesis' fourth paper, which develops a new approach for disaggregating U.S. annual building energy consumption to a sub-annual baseline in order to capture its hourly, seasonal, and regional variations and the resulting costs and emissions of these.

This new time-sensitive framework for building electricity consumption is essential for appropriately assessing the savings potential of energy efficiency and flexibility measures in buildings in order to inform investments in building technology R&D. The paper finds that the savings potential for these measures can be considerable (around \$2–8 billion and 5–26 $MtCO_2$ for summer cooling measures, alone). It also shows how a proper comparison of energy efficiency and flexibility measures is only feasible using a time-sensitive framework, which permits more holistic evaluations of different demand-side solutions. As energy efficiency and flexibility are the two cornerstones of electricity sector decarbonisation, understanding how they can best be deployed in tandem is increasingly important.

In summary, this thesis has several key conclusions for building energy research. First, the thesis shows how recent advances in statistical learning and the sub-fields of regularisation and sparsity are able to address numerous challenges inherent to building energy research. It shows how these methods can be used to identify key drivers of electricity consumption in complex building environments across varying geographical contexts. Second, while the thesis cannot claim to have investigated every potential factor that influences residential electricity demand, it has tested a large and diverse range of structural and occupant-related factors and has shown that although characteristics of the physical dwelling are important for understanding variations in consumption, so are aspects of occupant behaviour and time use, which are studied far less in the literature. Specific patterns of activities, especially those related to housework, recreation, and food, are shown to influence electricity consumption at different times of day. Evidence on these time-dependent relationships is lacking but is imperative for designing better demand reduction and response interventions in the residential sector. Third, the thesis has presented a new framework for time-sensitive valuation of energy efficiency and flexibility measures and has shown how such an approach can serve future energy researchers in their work to further characterise and quantify the potential of demand-side solutions for mitigating climate change.

Decarbonising the energy system at the speed and scale needed to limit global warming to 1.5°C is a challenge that will require much greater ambition across all sectors of the economy. Accelerating demand-side solutions to increase ambition for climate change mitigation has only recently become a focal point for the international energy and climate policy community. It is therefore of great importance, now more than ever before, that energy researchers deliver the analytical tools, data, and evidence necessary to realise these new pathways to meet formidable emissions reduction targets. This thesis aims to contribute to these needs for the residential building sector. The transition to a post-carbon society is imminent, but ensuring it takes place without economic, social, and environmental risks is not. Demand-side solutions that take into account the diverse preferences, attitudes, and behaviours of the people at the heart of the energy system can enable a just and equitable transition.

Appendix: Co-author statements

Papers 2 and 3 of this thesis were co-authored with Marina Diakonova and Philipp Grünewald from the University of Oxford Environmental Change Institute. I conducted the majority of the work, including the analysis and drafting of the manuscript. Marina and Phil provided guidance on the analysis as well as comments on various drafts. Marina and Phil have provided a declaration (below) which states that I, Aven Satre-Meloy, contributed the majority of the work on both of these papers.

Paper 4 of this thesis was co-authored with Jared Langevin from Lawrence Berkeley National Laboratory. I conducted the majority of the work, including the analysis and drafting of the manuscript. Jared provided guidance on the analysis and methods as well as comments on various drafts. Jared has provided a declaration (below) which states that I, Aven Satre-Meloy, contributed the majority of the work on this paper.

6 September 2019

As a co-author of the following papers that are being submitted as part of a student's Doctor of Philosophy thesis in the School of Geography and the Environment at the University of Oxford, I am writing to confirm that Aven Satre-Meloy undertook the majority of the work.

Name Marina Diakonova

Signature Diakonova

6 September 2019

As a co-author of the following papers that are being submitted as part of a student's Doctor of Philosophy thesis in the School of Geography and the Environment at the University of Oxford, I am writing to confirm that Aven Satre-Meloy undertook the majority of the work.

Name Philipp Grünewald

Signature 

6 September 2019

As a co-author of the following paper that is being submitted as part of a student's Doctor of Philosophy thesis in the School of Geography and the Environment at the University of Oxford, I am writing to confirm that Aven Satre-Meloy undertook the majority of the work.

Name Jared Langevin

Signature