Development of a profile-based electricity demand response estimation method for reducing uncertainty, as informed through a review of aggregator assessment processes and existing estimation methods

Submitted in fulfilment of the requirements for the degree of Doctor of Engineering (EngD)

UNIVERSITY OF READING
Technologies for Sustainable Built Environments Centre
School of Built Environment

Mitchell Curtis

February 2018

Abstract

This engineering doctorate (EngD) thesis has investigated and improved the understanding of Demand Side Response (DSR) aggregators, DSR estimation methods, and developed a new load profile-based estimation method. The primary motivation for this research was to develop and improve the understanding of different DSR estimation methods and their effectiveness for assessing new sites as suitable for DSR. DSR aggregators play a key role in facilitating DSR uptake by providing over 80% of DSR capacity. Therefore, this research has focused on the estimation methods that a UK-based aggregator uses to determine the suitability of new end users. As an intermediary in the DSR assessment and programme enrolment process, aggregators need to ensure that each end user site is suitable for DSR. Otherwise, both the aggregator and the end user could be negatively impacted if financial returns from participation fail to cover DSR implementation costs. Therefore, this research was undertaken with the aim of better quantifying the uncertainty in DSR estimation methods for new sites, with a view to improving the assessment of their suitability to participate.

The research was undertaken in conjunction with KiWi Power Limited, a UK-based DSR aggregator, by establishing and then addressing three interlinking objectives. The first objective mapped out the criteria used by KiWi Power to determine the suitability of an end user's site for DSR and found that the highest priority for KiWi Power during the assessment process is understanding the DSR potential of a site's assets. The second objective compared the outcome uncertainty and information input requirements of four existing DSR estimation methods using as the example asset HVAC Chillers and their sub-meter usage data from two UK hotel sites. The comparison results showed a range of uncertainty levels which produced mean average percentage error (MAPE) levels of between 39% to 159%, with the estimation methods costing between £10 to £180 to perform on new sites. The third objective developed and evaluated a new method that uses load profiles of assets to reduce the uncertainty of DSR potential estimation during an aggregator's assessment process. The new method compares favourably against the existing DSR estimation methods, as it generated the second lowest MAPE level of 46.5% with an estimated usage cost of £26. The new method demonstrated additional benefits of being usable earlier in the assessment process for a new site when compared to the existing methods, and offered the ability to use pre-calculated uncertainty levels enabling users to adjust the estimation outputs based on an organisation's risk appetite.

Declaration

I confirm that this is my own work and the use of all material from other sources has been properly and fully acknowledged.

Signed: Date:

Acknowledgments

There are many people to thank for helping me achieve this thesis. I am very thankful and lucky to have two great supervisors at the University of Reading: Jacopo Torriti and Stefan Smith. Both have been instrumental in providing knowledge, advice and the motivation needed to undertake this research. I thank everyone at KiWi Power who helped support my research. I thank Jenny Berger and Emma Hawkins at the University of Reading Technologies for Sustainable Built Environments (TSBE) Centre who have provided great support and made the centre a wonderful place to be a part of. I thank the Engineering and Physical Sciences Research Council who provided the research grant reference EP/G037787/1 enabling this thesis to be undertaken. I thank all my fellow TSBE researchers who have made the last four years enjoyable and intellectually exciting. I thank my family for their encouragement and support from the other side of the world. Finally, and most importantly, I thank Miriam Knight, who has supported me throughout this journey and without whom I could not have completed this thesis.

Table of Contents

List of Tables

List of Figures

List of Acronyms

AC	Air Conditioning
ADE	Association for Decentralised Energy
ANN	Artificial Neural Networks
BM	Balancing Mechanism
BMS	Build Management System
CCGT	Combined Cycle Gas Turbines
CHP	Combined Heat and Power
CI	Confidence Interval
CRM	Customer Relationship Management
CSV	Comma Separated Values
DNO	Distribution Network Operators
DR	Demand Response
DRQAT	Demand Response Quick Assessment Tool
DRRC	Demand Response Research Center (Lawrence Berkeley Laboratory)
DSM	Design Side Management
DSR	Demand Side Response
ENGD	Engineering Doctorate
FCDM	Frequency Control by Demand Management
FCM	Fuzzy C-Means
FFR-Primary	Firm Frequency Response Primary
FFR-Secondary	Firm Frequency Response Secondary
HVAC	Heating, Ventilation, and Air Conditioning
ICT	Information and Communications Technology
KOMP	KiWi Operation Management Platform
kW	Kilowatt
MAPE	Mean Average Percentage Error
MBE	Mean Bias Error
MW	Megawatt
Non-BM	Non-Balancing Mechanism
OCGT	Open Cycle Gas Turbines
OOS	Out Of Sample
PIP	KiWi Power's 'Power Information Pod'
SEDC	Smart Energy Demand Coalition
SOM	Self Organizing Maps
STOR	Short Term Operating Reserve
TSO	Transmission System Operators

1 Introduction

Demand Side Response (DSR) is increasingly seen as an important enabler of future low carbon electricity systems. DSR aggregators play a key role in enabling DSR by providing access to DSR programmes for end users who would otherwise be unable to meet mandatory participation requirements. For example, the Short Term Operating Reserve (STOR) programme requires participants to provide a minimum response of 3MWs. As an intermediary in the process, aggregators need to ensure that an end user’s site is suitable for DSR, otherwise the time and material costs of enabling an unsuitable site for DSR could outweigh the benefits to both parties of participation in DSR programmes. However, there is no existing research on how aggregators decide if a site is suitable for profitable DSR. There is also very limited research on methods used to estimate the DSR potential of a site, an essential element for calculating profitability. Therefore, this engineering doctorate (EngD) thesis focuses on increasing knowledge about how DSR aggregators work, the uncertainty levels of existing site level DSR estimation methods, and then developing a new model that uses asset usage profiles to reduce the uncertainty of DSR site estimation.

This chapter introduces the research being undertaken. Section 1.1 provides background on DSR and the importance it plays in future energy systems. Section 1.2 outlines the research context and motivation. Section 1.3 defines the research aim, objectives, and structure of the thesis. Section 1.4 concludes with a summary of key research contributions.

1.1 Background

The electricity supply sector is undergoing a paradigm shift that will result in the existing, unidirectional electricity supplier to consumer relationship model being replaced by an actively-managed network model, comprised of bidirectional flows of electricity and information from all stakeholders (Hadjsaid & Sabonnadiere, 2013). This new model is referred to as the 'Smart Grid' concept. It is driven by a number of factors which include energy markets' deregulation, increased distributed generation, security of supply and increased usage of intermittent renewable energy sources (Berger & Iniewski, 2012; Blumsack & Fernandez, 2012). One of the key enabling elements for the 'Smart Grid' concept is the change in the electricity supply balancing model from being unidirectional (supply meeting demand) to bidirectional, where demand is managed to meet supply ability. The definition for this new balancing approach is often contested as Warren (2014) outlines in his research on this area, concluding that the generally accepted term is "Demand Side Management" (DSM), defined as:

> *"Demand side management (DSM) refers to technologies, actions and programmes on the demand-side of energy meters that seek to manage or decrease energy consumption, in order to reduce total energy system expenditures or contribute to the achievement of policy objectives such as emissions reduction or balancing supply and demand."*

A sub-category of DSM consists of "Demand Side Response" (DSR) (also referred to as "Demand Response"). While there are many definitions of DSR, this research adopts the definition of the UK Association for Decentralised Energy (ADE, 2018) '*DSR is where energy users change their electricity consumption patterns in response to a signal or incentive to help balance the system*'. DSR initially became popular during the 1970s energy crisis (Nadel & Geller, 1996). While interest varied in the last decades of the 20th century, a renewed focus on DSR and its potential has been seen in the 21st century given its abilities to support current energy saving initiatives, improve the security of supply and enable new electricity generation forms (Postnote 452, 2014). Strbac (2008) further outlines the benefits of DSR as its capability to reduce expensive spare generation capacity, improve transmission efficiency and support generation from intermittent renewables. These benefits mean that DSR is increasingly seen as an important enabler of future low carbon electricity systems.

A report by the Smart Energy Demand Coalition (SEDC, 2015) highlights the importance of DSR uptake for meeting the EU 2030 Energy Strategy objective of renewable energy achieving at least a 27% share of overall energy consumption. DSR's capability for managing the demand side of the electricity system equation enables greater usage of variable renewable generation sources without having to rely on alternative costly measures for power generation, including backup generation and storage solutions (Barton et al., 2013). The UK transmission operator, The National Grid, sees DSR as a key enabler for managing future grid variability, and their Power Response programme aims to achieve 30-50% of balancing capability from DSR by 2020 (National Grid, 2017). Increasing support for the future usage of DSR is also backed by commercial and industrial end users. By way of example, a 2017 survey of UK businesses by The Energyst found that 77% of 96 respondents who are not yet using DSR would be interested in participating in DSR programmes in the future (The Energyst, 2017).

1.2 Research Context and Motivation

Participation of end users in DSR is primarily motivated by financial incentives. DSR market providers offer financial rewards when usage of DSR is cheaper than the traditional methods of managing network demand, which involves increasing and decreasing supply through usage of large electricity generation plants. During peak electricity usage times, for example, it can be cheaper to pay a DSR participant to reduce demand than it is to pay a generator to increase supply. Transmission System Operators (TSO) have the task of managing the supply of electricity between generation and distribution, which has resulted in TSOs traditionally stepping in as the main providers of DSR markets. TSOs are concerned about large-scale balancing of the electricity network (often at a country level), which results in their DSR programmes generally prescribing minimum MW participation requirements because their systems often rely upon having manual processes that only allow access for a small number of providers. These minimums result in many potential end users not being able to participate directly in DSR, and instead need to use a DSR aggregator.

Aggregators' importance for facilitating uptake of DSR is highlighted in a 2017 survey of UK businesses, with responses showing that 66% of DSR end users access the market via an aggregator (The Energyst, 2017). A similar pattern is seen internationally based on the Smart Energy Demand Coalition reporting which sets out that 80% of DSR capacity in the United States and other markets is provided by aggregators (SEDC, 2017). Aggregators are set to continue playing a key role in the UK DSR market based on the outcomes of the 2016/2017 Capacity Market auctions (Ofgem, 2017a). In the 2016 T-4 Auction all 1.4GW of DSR capacity was won by aggregators, and in the 2017 turndown DSR only auction aggregators won 82.7% of capacity (see section 2.1 for more detail on the Capacity Market and DSR products). The turndown DSR auction also highlights the reliance by aggregators on new sites for increasing uptake of DSR, as 90% of the tendered capacity for the 2017 auction was unproven, which means that new sites are required to fill this capacity. Consequentially, aggregators are reliant upon continuing to find new sites if they are to ensure adequate capacity for securing the awarded tenders. However, while aggregators play a key role in the DSR market and its uptake there is very limited research on how they assess the suitability of new sites for DSR. With increasing usage of DSR and recognition of its potential for enabling greater use of variable renewable generation sources, the role of aggregators will continue to increase. Therefore, it is important to understand how DSR aggregators determine the suitability of an end user and the challenges faced during the assessment process, to ensure DSR participation is maximised.

A crucial element in determining the suitability of a new site is understanding the potential financial returns and whether it is suitable for current DSR programmes. To gain this understanding requires estimating the site's potential, which is informed by the available electricity assets capable of being used for DSR. Despite the importance of the estimation process for determining site suitability for DSR, there is very limited research on DSR estimation methods with only three methods being referenced in existing research. The existing research outcomes also offer limited comparisons of methods to understand how they compare in terms of estimation error and informational requirements.

This lack of research on aggregators and DSR estimation methods is the primary motivation for this thesis. This knowledge gap was recognised by KiWi Power, a UK based DSR aggregator operating since 2009. To enable research into this area KiWi Power chose to be an industry sponsor of an Engineering Doctorate (EngD) in conjunction with the University of Reading. The EngD programme combines doctoral-level research in an industrially relevant field, with a taught element to develop technical training. In the case of this project, the initial problem of DSR estimation was proposed by KiWi Power and expanded on to fulfil the requirements of a doctoral research project.

1.3 Thesis Aim, Objectives, and Structure

The aim of this research is to better quantify uncertainty in DSR potential estimation for new sites, with a view to improving the assessment of their suitability to participate. This aim is achieved through the following interlinking research objectives:

1. To map out the criteria used by an aggregator to determine site suitability for DSR.

2. To perform a comparison of the outcome uncertainty in DSR potential estimation methods, evaluated against the level of informational requirements of those methods.

3. To develop and evaluate a model that uses asset usage profiles to reduce the uncertainty of DSR potential estimation during an aggregator's assessment process.

All three objectives provide key contributions towards addressing current gaps in the body of knowledge about DSR. The first objective reviews how DSR aggregators work and provides context to this thesis, increases knowledge about aggregators, and identifies the importance aggregators place on DSR estimation for determining a new site's suitability for DSR. The second objective moves the research focus to DSR estimation by reviewing and comparing current methods, providing new knowledge about how existing methods and their uncertainty levels and informational input requirements compare and contrast. The final objective develops a new DSR estimation method using load profiles, which when compared to the existing DSR estimation methods, improves uncertainty outcomes.

As an EngD project, the three objectives were undertaken as separate research strands that interlink to achieve the thesis aim. The thesis starts with a review of existing research on the subject matter and identified objectives. A chapter is then devoted to each objective and covers specific background information, research methodology, and results. The final chapter brings together the outcomes of each objective and draws out overall implications. Figure 1 shows a schematic representation of the thesis structure to illustrate how the chapters and content are organised.

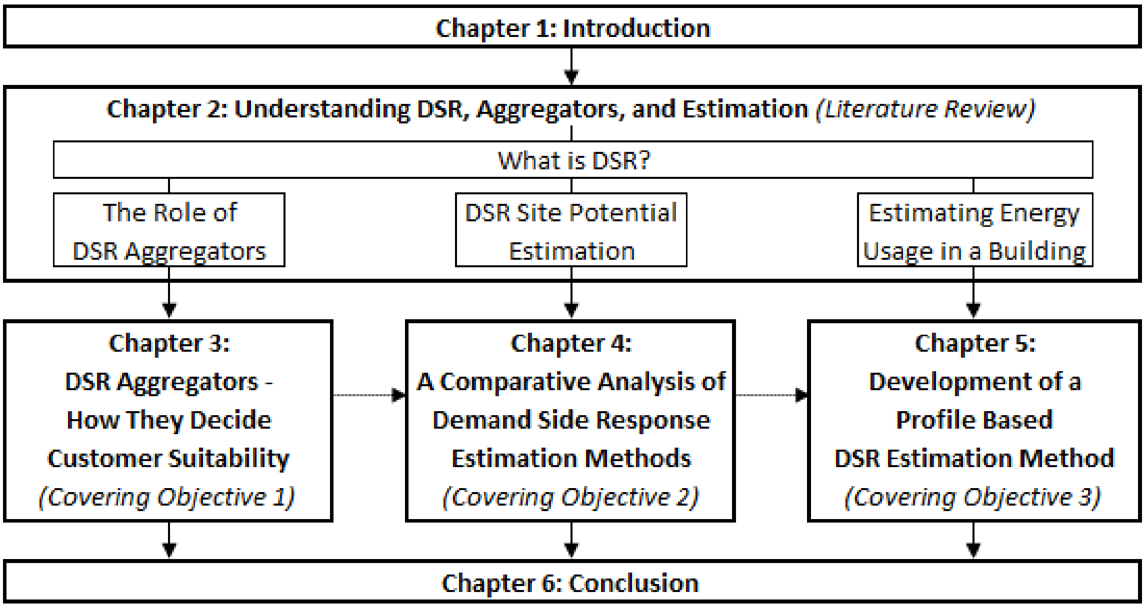

Figure 1 - Schematic Overview of Thesis Structure

The following briefly summarises each chapter:

- **Chapter 2 – Understanding DSR, Aggregators, and Estimation**

  This chapter reviews existing literature, starting by providing a general overview of DSR, including how it works, the different types, applications, and barriers to uptake. The review then focuses on summarising known information about each research objective and the knowledge gaps that are addressed by this research.

- **Chapter 3 – DSR Aggregators: How They Decide Customer Suitability**

  DSR aggregators play a key intermediary role in facilitating the uptake of DSR by enabling individual non-domestic end users to access complex system operator programmes. Chapter 2 identifies that there is limited research on how DSR aggregators determine the suitability of an end user site. Given this gap, this chapter addresses the first research objective by increasing the body of knowledge on aggregators by examining KiWi Power's acquisition process, and interviewing KiWi Power employees to identify key end user assessment tasks and the reasons why certain end user sites are deemed unsuitable. The results showed that the DSR estimation process is a key task when deciding the suitability of a new site.

- **Chapter 4 – A Comparative Analysis of Demand Side Response Estimation Methods**
  Chapter 2 identified that there is limited research on estimation methods and no studies that compare all known methods. This chapter addresses this gap and the second research objective by comparing four non-domestic DSR estimation methods to provide insights about uncertainty levels based on each method's input requirements. Each method is deployed to estimate the DSR potential of HVAC chiller assets at two UK hotels over two years. The results show the methods have a range of error levels from the highest Mean Average Percentage Error (MAPE) of 159% to the lowest MAPE of 39%.

- **Chapter 5 – Development of a Profile Based DSR Estimation Method**
  This chapter creates a new method that addresses the third research objective by developing and evaluating a model that uses asset load profiles to help reduce the uncertainty of DSR potential estimates. The method is developed using a seven-stage process: (1) data selection; (2) data preparation; (3) training and testing dataset selection; (4) profile creation; (5) profile evaluation; (6) method and input optimisation; (7) final profile creation and comparison. The results showed that when compared to the existing methods, the new profile method has the second lowest MAPE of 46.5%, and second lowest per usage cost of £26. The new method also provides additional benefits over the existing methods by providing an insight into uncertainty at the time of estimation which can then be used to adjust the estimate or inform business decisions based on the user's risk appetite.

- **Chapter 6 – Conclusion**
  Each research objective was addressed in its own chapter that provides objective specific results, discussion and conclusion sections. The concluding chapter reviews the outcomes for each objective before determining the overall implications of the research, including further potential research areas.

1.4 Key Contributions

This research has provided important contributions to the general body of knowledge on DSR and the commercial operations of KiWi Power. Table 1 provides a summary of the knowledge dissemination activities resulting from this research. The research has also contributed many benefits to KiWi Power throughout the research. Examples include: development of a web-based DSR estimation tool that has decreased analysis time and errors of new sites, while providing improved sales collateral that has increased uptake of new clients; creation of a winning £1.5 million Innovate UK grant; improving the DSR tendering process that resulted in KiWi Power being awarded the highest priced contracts from the National Grid.

Table 1 - External Contributions Made During Research

Output Title	Output Description	Chapter
SMEs - A New Market for Demand Side Response	Authored short paper and poster for Technologies for Sustainable Built Environments 2014 Conference, July 8, 2014	2
Provided evidence on '*Electricity Demand-Side Measures*'	Co-authored response to the Energy and Climate Change Committee's inquiry on electricity demand-side measures, August, 2014	2
Review of Barriers to Uptake of Demand Side Response in Medium Sized Businesses	Authored paper and presented at the 38th International Association of Energy Economics Conference, May 24-27, 2015	2
Barriers to Uptake of Demand Side Response in Medium Sized Businesses	Authored short paper and presented at Technologies for Sustainable Built Environments 2015 conference, July 7, 2015	2
Overview of the UK Demand Response Market	Presentation given at the EPFL Demand Response Workshop/Conference, September 11, 2015	2
Using building energy profiles for improving demand response services	Poster presented at the DEMAND Conference, April 13-15, 2016	5
Case Study on a UK Commercial Demand Side Response Implementation	Co-authored paper for the India Smart Grid Forum Conference, March 14-16, 2016	3
Improving Demand Side Response Prediction using Building Power Usage Profiles	Authored short paper and poster for Technologies for Sustainable Built Environments 2016 conference, July 5, 2016	5
Demand Side Response Opportunities and Policies in the UK	Presentation given to Chinese Energy Policy and Implementation delegation, October 10, 2016	2
Provided evidence on '*A Smart Flexible Energy System*'	Co-authored response to BEIS call for evidence on 'A Smart Flexible Energy System', January, 2017	2
Demand Side Response Aggregators: How do they decide customer suitability	Authored peer reviewed paper published in IEEE proceedings and presented at the 14th European Energy Market Conference, June 6-9, 2017	3
Comparison of Demand Response Estimation Methods	Authored short paper and presented at Technologies for Sustainable Built Environments 2017 conference, June 27, 2017	4
Demand Side Flexibility and Responsiveness: moving demand in time through technology	Lead author of peer-reviewed chapter in the book 'Demanding Energy: Space, Time and Change' published 2017, Palgrave Macmillan, Cham	2
A Comparative Analysis of Demand Side Response Estimation Methods in Buildings	Lead author of paper published in 'Energy and Buildings' journal on 30 July 2018	4

2 Understanding DSR, Aggregators, and Estimation

This chapter formulates an understanding of the existing body of knowledge relevant to the research aims and objectives. It serves two primary purposes: providing an overview of DSR, as subject matter referenced throughout the thesis; and identifying the knowledge gaps in DSR literature, which the research objectives then address. Section 2.1 provides an overview of DSR and explains the purpose of DSR, the providers, the different types, and the barriers to participation.

Sections 2.2 to 2.4 explain the outcomes from a literature review in areas relevant to each research objective. For the first objective '*To map out the criteria used by an aggregator to determine site suitability for DSR',* section 2.2 reviews the purpose of DSR aggregators, their role in the electricity system, and the issues faced. Section 2.3 then narrows the focus to DSR estimation methods for the second objective *'To perform a comparison of the outcome uncertainty in DSR potential estimation methods, evaluated against the level of informational requirements of those methods'.* This review identifies the reasons why DSR estimation is important for DSR uptake, and the existing methods available. Section 2.4 then addresses literature affecting the third objective *'To develop and evaluate a model that uses asset usage profiles to reduce the uncertainty of DSR potential estimation during an aggregator's assessment process'*. This review examines forecasting and modelling methods currently used to estimate energy usage in a building, which informs identifying potentially applicable approaches for developing a viable model for addressing the third objective. This chapter concludes with section 2.5, which provides a summary of the knowledge gaps found for each objective.

2.1 What is DSR?

As explained by Grünewald and Torriti (2013), DSR is seen as temporarily reducing the metered load of a site, enabling reduction in consumption to be achieved by either reducing usage at the site, hereby referred to as '**Turndown DSR**', or offsetting network supply with onsite generation, hereby referred to as '**Generator DSR**'. The motivational factors driving end user DSR participation are classified as either being implicit or explicit (SEDC, 2016). Implicit DSR covers price-based measures, whereby the end user may reduce demand based on the price of electricity. Explicit DSR covers incentive-driven measures, whereby the end user is requested directly to reduce demand based on an external signal. As 80% of DSR is currently provided by aggregators, who rely primarily on explicit DSR, means that this thesis focuses on explicit DSR (SEDC, 2017). The explicit DSR type uses Turndown DSR or Generator DSR techniques to temporarily reduce site demand based on incentive-driven measures from a DSR market provider. The sites are then financially compensated by the market provider for participating. The remainder of this section covers in detail the role of market providers and how Generator DSR and Turndown DSR is undertaken to form a DSR subject matter base that is referred to throughout this thesis.

2.1.1 DSR Market Providers

Based on a recent UK survey by The Energyst (2017), 79% of DSR end users are motivated by financial incentives, followed by corporate social responsibility at 26%. DSR market providers offer financial rewards when usage of DSR is cheaper than alternative methods of managing network demand. During peak electricity usage times, for example, it can be cheaper to pay a DSR participant to reduce demand than it is to pay a generator to increase supply. As TSOs have the task of managing the supply of electricity between generation and distributors, they have traditionally emerged as the main DSR market providers. This is reflected in the 2017 Smart Energy Demand Coalition report on DSR in Europe, which shows that almost all DSR products are provided by TSOs and only limited trials from Distribution Network Operators (DNO) or end user suppliers (SEDC, 2017). For the purposes of this research, the UK is used as the primary territory for examining the types of DSR markets provided with other country schemes being referenced when applicable, as this is the locale of this EngD research project and place of business for the aggregator studied.

The UK TSO, the National Grid, is the primary DSR market provider. The UK DNOs Western Power Distribution (WPD) and UK Power Networks (UKPN) have undertaken trials of DSR products to understand how they can be used to manage local network constraints. However, WPD and UKPN are each yet to implement commercial DSR programmes (UKPN, 2017; WPD, 2017). In contrast, the National Grid has long-running commercial programmes that can be used for DSR, like the Short Term Operating Reserve (STOR), a programme that has been operational since 2006. The National Grid DSR products form part of their overall Balancing Services, which aim to '*balance demand and supply and to ensure the security and quality of electricity supply across the GB Transmission System'* (National Grid, 2017d). For DSR, the range of products fall under two main categories: (i) Request Driven, and (ii) Frequency Response.

Request Driven DSR products are not automatic, and are only undertaken if requested by the National Grid. There are two products within this category, STOR and Capacity Market:

- **STOR** has been the UK's traditional method for providing DSR since it started in 2006. The National Grid uses this product to have approximately 3000-3500MW of '*reserve power in the form of either generation or demand reduction to be able to deal with actual demand being greater than forecast demand and/or plant unavailability'* (National Grid, 2017f). The split of generation and demand reduction capacity is not provided as this is reviewed in sections 2.1.2 and 2.1.3. To provide this product the National Grid uses a tendering process (run 3 times per year) to procure electricity reserve from companies (directly or via an aggregator) that can meet as requirements (National Grid, 2017f): providing at least 3MW of generation or demand reduction; deliver full capacity within 240 minutes of notification (though the majority of tenders are sub 20 minutes (National Grid, 2017d)); provide capacity for at least 2 hours; be available for 3860 hours a year during a morning and afternoon window. The product is funded by paying winners of the tenders an availability price (£ per MW per hour) and utilisation price if called on (£ per MW per hour).

- **Capacity Market** was introduced in December 2014 as a part of the Electricity Market Reform programme, which is charged with ensuring the future security of the UK electricity supply by making sure that sufficient capacity is able to be provided during times of system stress (Ofgem, 2017a). The Capacity Market is technology neutral and therefore accepts capacity from generation, DSR and interconnectors. The capacity is awarded through yearly auctions covering future delivery up to 4 years in the future with the 2016 T-4 auction acquiring 52,400MW, of which 1,411MW was from DSR Generator and Turndown sources (Ofgem, 2017a). There was also a special Turndown DSR only auction in March 2017 which resulted in 312MW of capacity being awarded. The participation requirements for the capacity market are: providing at least 500kW of capacity (directly or via an aggregator), deliver capacity 4 hours after being notified, provide capacity for at least 4 hours, and be available to deliver the capacity at any time during the operational year. The product is funded by paying winners of the auction a set yearly payment (£ per kW).

Frequency response DSR is automatically triggered based on changes in the network frequency level. Three products within this category are used for DSR: Firm Frequency Response Primary, Firm Frequency Response Secondary, and Frequency Control by Demand Management (FCDM):

- **Firm Frequency Response Primary (FFR-Primary) and Firm Frequency Response Secondary (FFR-Secondary)** trigger reduction based a predefined network frequency value that is normally between 49.5 and 49.7Hz (National Grid, 2017a). The trigger Hz varies across providers to prevent all capacity triggering at the same time. FFR-Primary needs to respond within 10 seconds of the trigger while FFR-Secondary permits 30 seconds for a response. FFR-Primary duration is short, at only 1 or 2 minutes, to give time for FFR-Secondary to engage, which is then required for up to 30 minutes. Both products have a minimum 1MW requirement (directly or via an aggregator) that is procured during a monthly auction. They are awarded through a tender process with winners being paid an availability fee (£ per MW per hour).

- **Frequency Control by Demand Management (FCDM)** is used to manage large deviations in frequency which can be caused by, for example, the loss of significantly large generation (National Grid, 2017b). It was originally setup to enable demand reduction of large electricity users like steel smelters, and will be stopped at the end of 2017 with users being converted to FFR products.

A summary of the DSR product requirements is provided in Table 2. The timeframes of each product have been plotted in Figure 2 to help demonstrate how they are similar yet slightly different. As can be seen in the plot, the fast-automated frequency based products are designed to rapidly respond to abnormal imbalances in the network and provide reduction for up to 30 minutes. This can occur around 10 to 30 times per year when unexpected events happen, like an unplanned disconnection of a large generator due to a fault (Grid, 2017). The 30-minute duration allows the National Grid enough time to decide on next steps. The next step would normally be to request a large generation plant to provide more capacity or, if that is not available, then a request to STOR providers to reduce demand while slower starting generators are brought online. Once the STOR capacity usage had been activated, the National Grid aims to balance the system within 2 hours using traditional large generators. The Capacity Market provides the National Grid with an extra level of security when a 'stress event' is likely to occur and available generation will not meet forecast demand. This is expected to be a rare occurrence, and has only been used twice in its first year of operation between 1st October 2016 and 30th September 2017 (National Grid, 2017c).

Table 2 - National Grid DSR Product Requirements

Programme	Response time	Max Duration	Minimum MWs	Trigger
FCDM	Within 2 seconds	30 minutes	3	Static Frequency Point
FFR – Primary	Within 10 seconds	1 to 2 minutes	1	Static Frequency Point
FFR – Secondary	Within 30 seconds	30 minutes	1	Static Frequency Point
STOR	Within 20 minutes	2 hours	3	National Grid Request
Capacity Market	After 4 hours	4 hours	0.5	National Grid Request

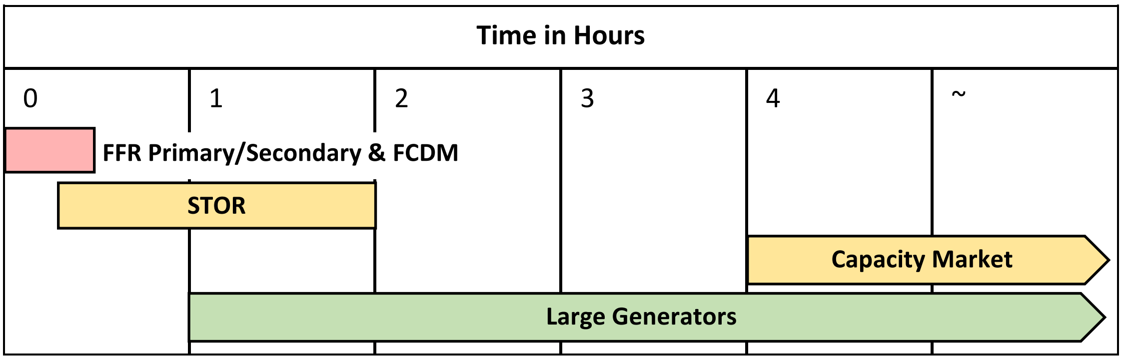

Figure 2 - National Grid Balancing Services

2.1.2 Generator DSR

Generator DSR refers to the usage of onsite generation sources as a means for temporarily reducing a site's electricity demand on the network supply. Traditionally this is done by using standby/backup generators. Andrews (2007) explains that backup generators are a logical choice for helping manage network variability as these are already designed to temporarily covering a site's demand in case of network supply failure. He also notes that generators need to be tested at load for 1 or 2 hours per month so participation in DSR provides the benefits of: covering the test running costs, generating income from an otherwise expensive and underutilised asset, and providing a useful service to the network. The suitability of generators to provide DSR is reflected in a review of a UK DSR aggregator which found that 84% of participants' DSR capacity was from sites with backup generators (Grünewald & Torriti, 2013).

Usage of generators for DSR is not without controversy, with a report by the Energy and Climate Change Committee recommending that '*the definition of demand-side response should exclude consumers turning on their own generation assets such as diesel generators. This agreed definition should be consistently and immediately applied by DECC, Ofgem and National Grid*' (Commons Select Committees, 2015). The Committee took a view that using generators does not comprise genuine demand reduction, and instead amounts to localised generation. One of the concerns with using diesel backup generators for DSR is the extra pollution compared to using traditional large-scale generation plants. A study by Lau & Livina (2015) compared the carbon emissions caused by different generator types used in the UK National Grid's STOR programme. The authors found that diesel generators had lower total carbon emissions compared to using gas turbines due to the ability of diesel to provide required electricity within minutes, conversely the turbines required longer to start-up and stop. Therefore, while diesel generators produce more emissions than turbines when running, this is offset by the longer operating times of turbines when used for STOR. However, backup generators are generally located in populated areas, and a study by Tong & Zhang (2015) showed that running these generators for DSR in areas with tall upwind or downwind building forms negatively impacts local air quality.

The amount of Generator based DSR in the UK is difficult to quantify due to a lack of published figures. The closest estimate can be obtained from a National Grid report on fuel usage by STOR programme participants (National Grid, 2015). The report outlined that there was 3,444MW of STOR capacity provided from generation and demand reduction sources. This capacity is split into two major groups, with the first group covering 1486MW (43%) of Balancing Mechanism (BM) sources which the report deems as generation sources that are '*connected directly to the GB Transmission System or are large enough to have to register in the Balancing Mechanism (BM)*'. The second group covers 1958MW (57%) of Non-Balancing Mechanism (Non-BM) sources, which the report deemed as being '*participants that are typically represented by smaller providers connected to the lower voltage distribution networks. When these operate, National Grid sees their impact on the transmission system as a reduction in demand and for this reason these providers – whether generation or load reduction services – can be referred to as "demand side" providers*'.

Figure 3 shows a breakdown of the 1958MW of Non-BM sources that the National Grid categorises as demand side providers. However, there is insufficient information to determine exactly how much DSR is using generators to replace site demand, and how much demand reduction is occurring due to distributed generation feeding into the local network. This report's values can be refined by removing sources not associated with site backup generators, namely the 219MW of Combined Cycle Gas Turbines (CCGT), 105MW of Combined Heat and Power (CHP), 151MW Hydro, and 368MW Open Cycle Gas Turbines (OCGT) which are either used solely for generating power for the local network, or are used as the primary electricity source for the site and therefore unlikely to lower site demand (and so not deemed as generator-based DSR). The 237MW Load Reduction category can be excluded too, if deemed as Turn Down DSR, i.e. DSR that does not use generators and is instead covered in the next section 2.1.3. Excluding these values leaves up to 878MW of potential generator-based DSR in the UK. Yet this amount is still not accurate as the remaining sources, namely diesel generators, are highly likely to contain some sites that are dedicated 'Diesel Farms', which only export to the network, and other sites that also export spare generation capacity. However, as this calculation is only aimed at providing a general idea of the generator-based DSR capacity in the UK, 878MW will be used as a reference point, which is a small amount of the estimated 20GW of backup generator capacity across the UK (Andrews, 2007).

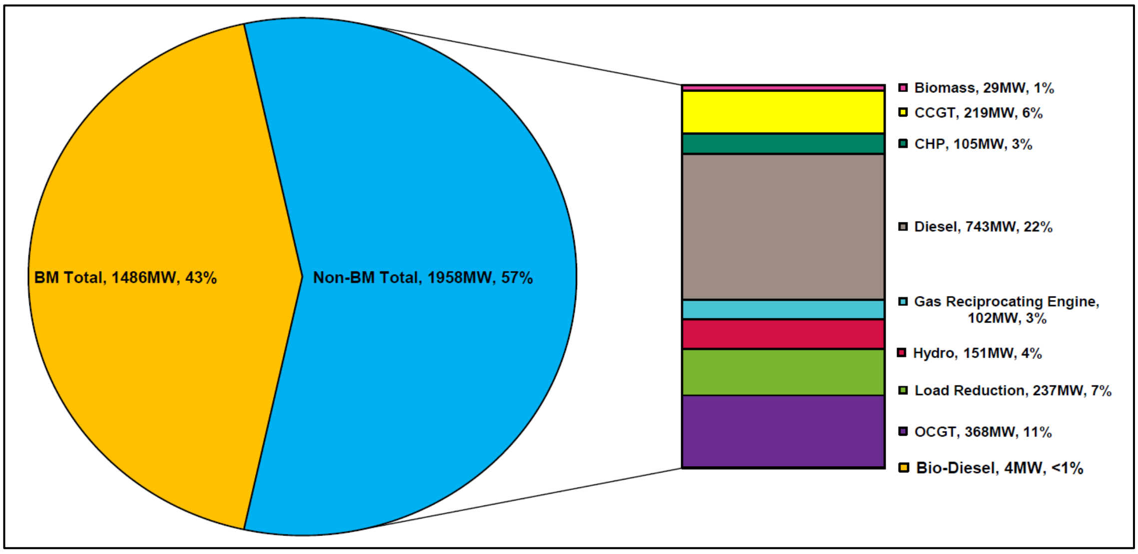

Figure 3 - Detailed Breakdown of STOR Non-Balancing Mechanism Sources.
Source: (National Grid, 2015)

An important consideration before using backup generators for DSR is determining whether there is any ability to export to the local network spare capacity which is not being used by the site. Capability to export spare capacity can improve a site's DSR suitability as the site owner will be able to claim any exported electricity as additional demand reduction. It also means that the generator can run at full capacity and not vary during operation, even for site-specific demand fluctuations. This has the effect of making it easier to forecast the site's DSR potential while increasing its revenue potential.

The ability to export capacity requires a G59 Mains Protection Relay connection (often just referred to as a G59 connection) and an export agreement. A G59 connection is necessary if the site plans on running its generator at the same time as also using grid electricity (Shenton Group, 2013). The G59 relay acts as a safety device, disconnecting the site from the grid in the event of a main's electricity failure. This disconnection prevents electricity being fed back onto the grid from the generator, which could otherwise endanger anyone fixing a local electricity network issue. If the site does not have a G59 connection but wants to utilise their generator for backup or for DSR, then the site must go into 'island mode', by operating the generator in isolation from the electricity grid and only providing enough electricity as required by the site (Clarke Energy, 2016). From a DSR point of view, the island mode is an acceptable method as the site's load is removed from the national electricity system and therefore reduces demand as required.

While obtaining a G59 connector to allow parallel running is straightforward (at a cost) the ability to also export can be difficult. This is due to the export constraints that apply to DNOs on selected parts of the network (Hoare, 2015). When the DNO has too much locally generated electricity enabling any more capacity will require upgrading the network which is normally cost prohibitive. As a result, a G59 connection with an export agreement is the preferred situation for generator-based DSR, yet can be difficult and time-consuming to obtain.

2.1.3 Turndown DSR

There are two main types of Turndown DSR, 'load shifting' and 'load reduction'. Load shifting achieves reduction in electricity demand during the requested period by moving the load to a different time. This approach is often used for heating and cooling services that can handle short stoppages without impacting users, and in situations where temporarily stopping production will not have a major impact to the user's business (Qureshi, Nair, & Farid, 2011; Wang & Li, 2013). In contrast, load reduction reduces usage that will not need to be recovered at a later time, for example temporarily turning off lights (Siano, 2014). The method used will depend on each end user situation, and is likely to be influenced by the financial benefits offered by the operator for end users receptive to applying methods to reduce their energy usage.

The amount of Turndown DSR in the UK is hard to accurately quantify due to limited published figures. However, a rough estimate can be determined from the figures that have been published. The first information source is the previously referenced National Grid report on STOR programme participants' fuel usage (National Grid, 2015). Figure 3's breakdown of fuel usage shows that 237MW (7% of overall STOR capacity) is from load reduction (which in this case is likely to also include load shifting DSR, as the National Grid uses the term 'load response' to cover both types in other non-fuel usage reports in 2017 (National Grid, 2017e)). The second source of information is the Turndown DSR only Capacity Market auction in March 2017, which resulted in Ofgem awarding 300MW (Ofgem, 2017a). These two MW figures cannot be combined as each could represent the same turndown sources. There are no known sources either of information on the other sources of Turndown DSR like Frequency Response. Therefore, based on the figures obtained from published reports and assuming additional amounts have been provided via other sources, a rough estimation of Turndown DSR in the UK could comprise between 300-500MW. A report by Element Energy (2012) further illustrates the potential of Turndown DSR for the UK, as this report determined three potential non-domestic Turndown DSR capacity levels in the UK: Conservative 1.2GW, Moderate 2.5GW, and Stretch 4.4GW. This variation highlights the opportunities yet clear uncertainty in the DSR potential offered by the electrical assets in non-domestic businesses.

Understanding the Turndown DSR potential of electrical assets is difficult as suitability will depend on the assets' characteristics and also DSR programme requirements. One way of determining suitability is to assess the flexibility and responsiveness of appliances in relation to DSR. In a chapter on 'Demand Side Response: moving demand in time through technology' (Curtis, Torriti, & Smith, 2018) the concept of DSR flexibility and responsiveness is discussed in detail, applying the following definitions for 'flexibility' and 'responsiveness':

> ***Flexibility*** *- a measure of how rhythms of demand change in time. This can be operationalised as a technical measure of how patterns of the asset/device 'flex' or change in time and will vary based on several factors including: time of day, duration of change, intensity of change, and recovery time.*

> ***Responsiveness*** *- a measure of how fast the asset can respond to an event (change) request, given an economic or technical input (such as a peak or price signal).*

This thesis focuses on DSR estimation methods which are primarily concerned with determining the flexibility of appliances. Responsiveness of the appliances remains relevant, as this is used to then decide which DSR programme can be used. To understand how these two factors influence the suitability of appliances for DSR the seven non-domestic, non-industrial electricity consumption end use areas used by the Element Energy (2012) report will be reviewed. The seven end use areas are shown in Figure 4, including catering and computing which the authors had excluded as they deemed them as unsuitable due to lack of flexibility. To understand why end use areas and their associated electrical assets might or might not be suitable for DSR, each usage area will be reviewed using both the flexibility and responsiveness definitions.

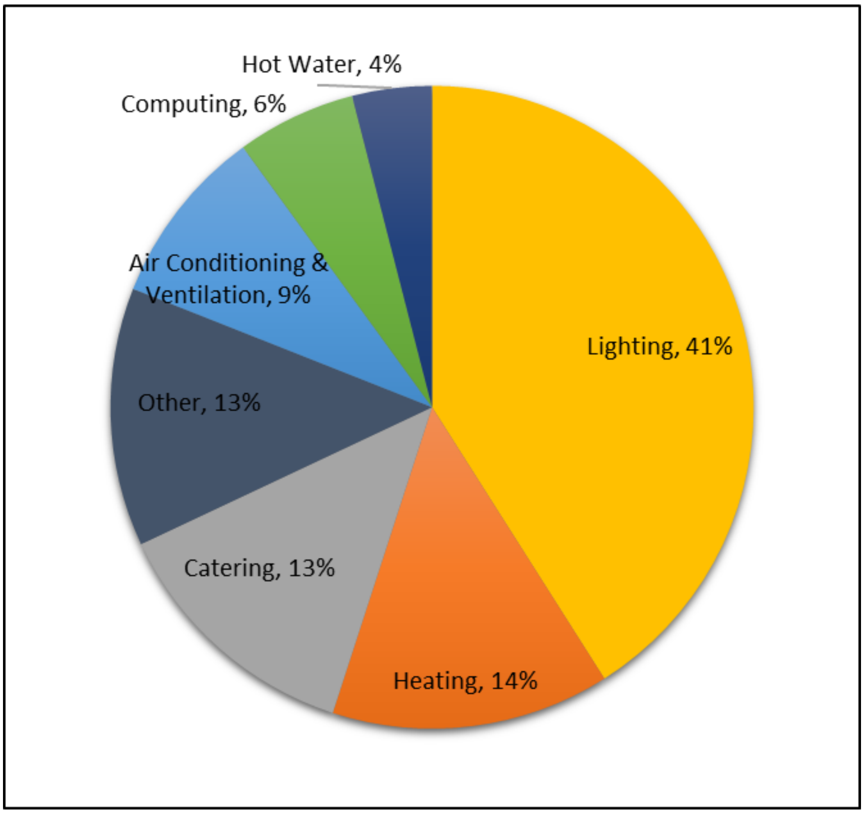

Figure 4 - 2009 UK Annual Non-Domestic, Non-Industrial Electricity Consumption by End-use. Source: (Element Energy, 2012)

Lighting has a very fast responsive level as it can be turned off or down almost instantly. However, its level of flexibility is subject to debate. Siano (2014) states that lighting is a proven solution for DSR with a wide range of applications (though no data or references are provided). A trial of DSR strategies in commercial buildings by Page et al. (2011) found dimming worked in an office environment with users not recognising the difference between DSR event and normal operation. A report by Lockheed Martin Aspen (2006) on 'Demand Response Enabling Technologies For

Small-Medium Businesses' found lighting to be an option, but only if the controls were already in place to allow for dimming or selective turning off. If the controls are not already in place, then the expense to install them would outweigh any benefits of DSR participation. The high level of electricity usage associated with lighting (41% as per Figure 4) means that it has been subject to recent energy efficiency drives including in the USA, which has resulted in office buildings reducing lighting's percentage of usage from 38% to 17% between 2003 to 2012 (EIA, 2017). This reduction reduces the potential of lighting appliances for DSR, expressly on the basis that if lighting can be turned off, then it probably should not be on in the first place. Therefore the flexibility of lighting is deemed medium to low which is supported by the Element Energy (2012) report which recognises that while lighting could offer a significant amount of DSR potential (over 50% of the quoted values are represented by lighting), there are a number of issues with its practical application due to costs of controls, for example, and the lack of scenarios where dimming or turning off lights would actually work in a commercial environment.

Heating in the form of electricity-based space heating has a fast response rate as most non-industrial heating systems can be turned off or down almost instantly. Its flexibility is generally regarded as being high to medium as it can be turned off or down for short to medium periods of time without impacting users, and it generally follows predictable patterns of usage (on during most of the day in winter, off in summer). This variance if flexibility is due to researchers having differing opinions on the period that heating can be turned off or down, without having a noticeable impact. Element Energy (2012) state that it can only be adjusted for a short period of approximately 15 minutes before impacting users. The Lockheed Martin Aspen (2006) report includes space heating control in most scenarios yet does not specify durations. Research on phase change materials by Qureshi et al. (2011) found that their usage can allow for heat load shifting of 2 hours or more. This means that using space heating for DSR will depend on each building's characteristics.

Catering has a fast to medium response rate as most cooking systems can be turned off or down instantly, conversely catering refrigeration systems can take a short amount of time to turned off correctly due to their usage of compressors. The flexibility of catering is deemed low due to the inability to shift cooking times based on needing to meet customer expectations (Element Energy, 2012; Grein & Pehnt, 2011a). One potential area of flexibility though resides in the thermal stores present in refrigeration systems, which can account for over a third of commercial kitchen electricity demand (Mudie, Essah, Grandison, & Felgate, 2014). However, using refrigerators for

DSR would require additional controls to ensure health and safety standards are maintained, and these are likely to make it economically unfeasible.

Computing has a low response rate as these devices need to be shut down in a controlled manner. Their flexibility is also low due to normally being required whenever turned on. Only the Element Energy (2012) report addresses using computers for DSR, and deems them to have a low potential due to the impact on users.

Air Conditioning & Ventilation has a medium response rate as research on hotel HVAC Chiller units by Curtis, Torriti, & Smith (2018) found that they can take on average 7.5 minutes to turn off as these appliances require sufficient time to slow down internal pumps to prevent components being damaged from freezing or blockages. The HVAC chiller units and ventilation system in general do offer high flexibility due to inertia in these systems allowing for delayed impacts to users. There is limited research on the amount of time HVAC systems can be turned off without impacting users. However, ad hoc studies generally indicate that air conditioning can be turned off for approximately an hour before users start to notice (Barton et al., 2013; Xue et al., 2014). Research on DSR assets often highlights HVAC as a being very suitable due to its limited short-term impact on users when turned off or down (Page et al., 2011). Element Energy (2012) states that HVAC can be interrupted for periods of up to 1 hour without impact to users, though without providing any evidence to support this claim. The Lockheed Martin Aspen (2006) report shows that most HVAC systems have the necessary controls for DSR, but that the systems are all generally proprietary. Therefore, each system needs to be addressed separately when assessing suitability and enabling DSR, and this can increase costs.

Hot Water has a fast response time as the heating elements can be turned off or down instantly without impact. Hot water flexibility is high when using storage tanks but low when using instant hot water systems. Sake et al. (2013) concluded that hot water tank systems are a viable source of DSR that can be implemented with little if any impact on users. Element Energy (2012) also claim that hot water tanks have a high DSR potential, yet highlight that instant hot water systems are unsuitable unless used in non-critical applications where cold water is sufficient (i.e. handwashing).

In addition to these seven commercial asset areas analysed, the UK National Grid guide to DSR also identifies the industrial specific areas of Pumps/Compressors and Industrial/Manufacturing

Processes (Grid, 2017). Each of these additional areas will be analysed to understand their potential for DSR.

Pumps/Compressors have medium to low responsive levels due to the need for controlled shutdown to prevent system issues. These appliances offer a range of flexibility levels depending on their application. Pumps offer medium to high flexibility when used in large-scale applications like waterworks, which are able to rely on built-in storage to allow pumps to be turned off for short periods of time (Menke at al., 2016). Compressors used in refrigeration/cold stores situations offer medium flexibility based on Grein & Pehnt (2011b) determining that small-scale refrigeration (retail size) could manage being turned off for less than half-hour and larger warehouse cold stores could also tolerate being turned off for up to an hour or longer, if able to lower cooling temperatures before an event.

Industrial/Manufacturing Processes have low response times for DSR due to any interruptions to the processes normally requiring sufficient notice to allow for a controlled stop. The flexibility will vary by industry category, yet these processes are generally deemed to have a medium to low level due to the likely impact on normal operations. A report by Agora Energiewende (2013) into German industrial DSR potential found that HVAC had the greatest potential followed by compressed air. This report noted that the energy-intensive production areas of cement, paper, chlorine, and steel can provide a reduction of up to 2 hours with sufficient financial incentive, yet required notifications of between 1 to 2 hours before the event.

Turndown DSR is comparatively more complicated than Generator DSR due to each asset having different levels of flexibility and responsiveness. Generator DSR is simpler as this technique is able to replace site demand without impacting normal operations, conversely Turndown DSR normally involves some form of impact to the site or its occupants. The range of flexibility and responsiveness levels across the asset categories can make it difficult to understand which areas are most suitable for Turndown DSR. Therefore, to help inform an understanding of the relationship between flexibility and responsiveness Figure 5 shows a phase space schematic of the Turndown DSR categories. The schematic shows that Hot Water Tanks have the greatest potential as these appliances are able to be instantly turned off, and have a high flexibility due to a low impact on users as the hot water stored in the tank is still accessible, even during an event. In contrast, computers are deemed to have the lowest potential as these devices take time to turn off and a low flexibility as turndown directly impact users by stopping them from using the computer. The AC & Ventilation area is highlighted on the schematic as this will be the primary focus area for this thesis based on available data sources. These different levels of flexibility and responsiveness result in barriers to uptake of DSR, as reviewed in the next section 2.1.4.

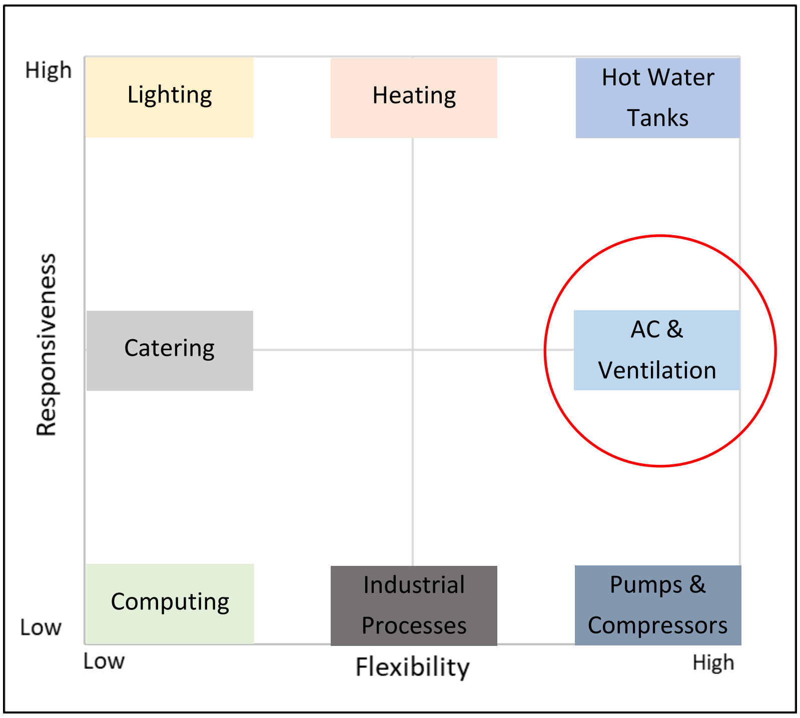

Figure 5 - Flexibility versus Responsiveness of Turndown Assets. The Red Circle Highlights the Focus Area of this Thesis.

2.1.4 Barriers to DSR uptake

Although DSR is an important element of future electricity systems, there are still many barriers preventing end user uptake. The research literature on DSR identifies many different types of barriers, which can make it difficult to understand how barriers interrelate and ultimately influence DSR uptake. Therefore, to facilitate understanding about the barriers and their potential for affecting DSR uptake Table 3 provides a matrix of the primary barriers raised by DSR relevant literature. Based on the literature, the barriers fit into the four main categories of: end user, regulatory, technical, and market. Many of the barriers in the literature could be interrelated, for example the cost of enablement barrier is linked to insufficient financial incentives. However, in this matrix, barriers are categorised based on the primary concern generated by a barrier. By way of example, enablement costs are primarily a technical issue, as the costs of enablement could be reduced through improved designs and economies of scale, and are not automatically a consequence of insufficient financial incentives. This section reviews each barrier category identified in this matrix to gain an understanding of the specific issues faced within each category.

Table 3 - Matrix of DSR Barriers Raised in Literature

Source Key:

1 = (Strbac, 2008) Demand Side Management: Benefits and Challenges

2 = (Owen, Ward, & Pooley, 2012) What Demand Side Services Could Customers Offer?

3 = (Cappers, MacDonald, Goldman, & Ma, 2013) An Assessment of Market and Policy Barriers for Demand Response Providing Ancillary Services in U.S. Electricity Markets

4 = (Warren, 2014) A Review of Demand-Side Management Policy in the UK

5 = (Nolan & O'Malley, 2015) Challenges and Barriers to Demand Response Deployment and Evaluation

6 = (Olsthoorn, Schleich, & Klobasa, 2015) Barriers to Electricity Load Shift in Companies: A Survey-based Exploration of the End User Perspective

7 = (SEDC, 2017) Explicit Demand Response in Europe: Mapping the Markets 2017

8 = (The Energyst, 2017) Demand-side Response: Shifting the Balance of Power: 2017 Report

Barrier Category	**Barrier**	**Barrier Research Source**							
		1	**2**	**3**	**4**	**5**	**6**	**7**	**8**
End user	Lack of DSR awareness/understanding	●							●
	Impact Concerns	●	●			●	●		●
	Risk aversion/trust issues		●				●		●
Regulatory	Regulations unfavourable for DSR	●		●	●	●	●	●	●
	Current regulations preventing DSR							●	
Technical	Lack of ICT infrastructure	●				●			
	Cost of enablement			●			●		●
	Equipment not suitable for DSR			●			●		●
Market	Lack of DSR market options	●		●		●		●	
	Insufficient financial incentives		●	●		●	●		●
	Traditional large generation bias			●	●	●			

2.1.4.1 End User Barriers

This category of barriers focuses on issues that end users have direct influence over, as opposed to the other barriers like financial incentives, for example, that while impacting a user are actually controlled by an external party. Table 3 shows that the most commonly raised end user barrier comprises concern about the impact of using DSR on business operations. This concern is exposed by end user surveys, with a survey of 78 UK businesses by The Energyst (2017) finding that business impact concerns ranked as the 3rd highest reason for not using DSR. In a survey about DSR barriers of 287 German industrial companies by Olsthoorn et al. (2015), the researchers found that business impact concerns ranked as the highest reason for not participating in DSR. The Olsthoorn et al. survey found that the industrial companies were most concerned about DSR participation stopping or interrupting operational processes that could impact the quality of a company's outputs and its ability to meet quality and delivery targets. Nolan & O'Malley (2015) also note that impact concerns can also occur for users who are already participating in DSR if they are inconvenienced while participating, causing them to disengage or demand higher incentives.

Another common end user barrier theme was lack of trust and risk aversion. Owen et al. (2012) point out that some businesses, like food manufacturing, have processes that are safety-critical, and therefore have no interest in taking unnecessary risks with their energy supply or operations in order to provide DSR. This could explain why the highest reason for not using DSR in The Energyst (2017) survey was due to equipment or processes not being suitable. The Energyst survey also highlights a lack of trust with allowing any third party to have control of a potential participant's systems, though this appeared to be less of a concern than other barriers, as handing control of equipment to others is only the second lowest reason for non-participation.

End user lack of awareness or understanding of DSR as a barrier preventing uptake was only raised twice in the selected literature. However, it does appear to be an ongoing issue as the first appearance of this as a specific barrier was in the 2008 article by Strbac (2008) and then again in the latest 2017 report by The Energyst (2017). The Energyst report found that 24% of survey respondents were not aware of DSR opportunities, while 18% did not understand enough about the market and options to make a decision. This last point reflects a concern raised by Strbac in 2008, specifically that there is a lack of clarity on DSR benefits that impact the business case for DSR.

2.1.4.2 Regulatory Barriers

The impact of regulation and policy on DSR uptake varies by country as highlighted in a report by the Smart Energy Demand Coalition (SEDC, 2017) that found the regulatory framework in Europe for DSR as improving, but still fragmented with major barriers. One specific barrier noted by the SEDC is that the majority of EU member states do not yet acknowledge the role of independent DSR aggregators in enabling uptake. This is preventing wider uptake of DSR as access is restricted to large end users. In the UK the level of government policy and regulation on the energy industry has varied over time with an increase in rhetoric about policy and regulation more recently due to the UK Government's climate change commitments, anxiety about securing future energy supplies and the public perception of energy companies as 'profiteering' (Stern, 2014). However, within the energy industry, the DSR market has only recently started to appear on governmental and regulators' radars, which suggests positive steps towards increasing its acceptance and usage or at least increasing awareness about the relevance of DSR to sustainable energy supply. The UK Government recognises the importance of DSR, yet in a review of UK DSR policy by Warren in 2014, he identified that they have not yet been able to provide the clear policy statements needed to provide the foundations required to increase DSR usage.

Progress on DSR policy has improved since 2014 with the introduction of the Capacity Market that '*provides a regular retainer payment to reliable forms of capacity (both demand and supply side), in return for such capacity being available when the system is tight*' (DECC, 2013). This policy specifically includes DSR as a measure to meet the Capacity Market mechanism's aims with the latest auction 2016 T-4 auction acquiring 52,400MW, of which 1,411MW was from DSR Generator and Turndown sources (Ofgem, 2017a). There was also a special Turndown DSR only auction in March 2017 which resulted in 312MW of capacity being awarded.

While regulatory barriers are being addressed in the UK they are still attributed as one reason for preventing greater DSR uptake based on the SEDC report. Its authors state that in the UK "*the opportunity for Demand Response is in principle higher than ever. However, due to poor policy development and design choices, that opportunity has not yet been realised*". The Energyst (2017) report also highlights the issue of continual regulation uncertainty and its impact on future investments, expressly for long-term battery projects. The Energyst does though see progress is being made with potential plans for allowing DSR aggregators to access the balancing market and wholesale market, which will improve their ability to be competitive.

2.1.4.3 Technical Barriers

As DSR is technical, and involves manually or automatically controlling electrical equipment via external signals, a range of technical barriers can prevent participation. One of the major technical barriers is end user equipment being deemed as unsuitable for DSR. The Energyst (2017) survey found unsuitability of equipment as the highest ranking reason for not participating, while the Olsthoorn et al. (2015) survey found it to be a medium ranked reason. The actual reason why equipment is deemed unsuitable is not specifically covered in the reviewed literature, though Cappers (2013) point out that the inability to meet DSR performance requirements (for example, responding to a DSR signal in the required timeframe) could deem equipment as unsuitable. The reasons could also be interlinked with business impact concerns, as previously noted in section 2.1.4.1, with users deeming a process to be uninterruptable due to safety reasons, which might relate to the technical ability to manage production conditions. For example, a large storage fridge could technically be turned off short amounts of time without impact. However, if the DSR system is unable to interact with the fridge's temperature monitor system and stop the DSR event if the temperature rises above an acceptable level, then the equipment could be deemed as technically unsuitable.

Equipment unsuitability could also be interlinked with the technical barrier of cost enablement. The cost of enablement barrier could be argued as comprising a financial barrier on the basis, for example, that if there were sufficient financial returns then these would justify the cost of participation. However, assuming that financial returns will not increase, then overcoming this barrier will require finding lower cost technical solutions. An example of this barrier is raised by Dr Oakes in The Energyst (2017) report "*There are a lot of pieces of kit out there that will read frequency arguably to National Grid's standards that cost £5-£10 a piece. But the only frequency meters that National Grid approves cost £300-£400. Putting that into a small site means the first couple of years benefit is gone and it is no longer economic.*"

The monitoring equipment required for DSR relates to the last technical barrier, namely the lack of Information and Communications Technology (ICT) readily available to support DSR. Strbac (2008) deems this to be a major challenge for DSR as he states that to support its implementation will require a significant deployment of the necessary advanced measurement and control devices. This challenge is reinforced by Nolan & O'Malley (2015) who point out that the currently installed electricity monitoring meters are not suitable for DSR event monitoring as these meters do not provide the required level of data detail needed to verify event outcomes.

2.1.4.4 Market Barriers

The balancing of participation costs against the financial returns offered from available DSR market options influences the capability of end users to participate. A lack of DSR market options compounds this issue, as highlighted by the UK who are considered as having one of the most advanced DSR markets in the EU yet still do not have a specific market programme for DSR (SEDC, 2017; Torriti & Grunewald, 2014). DSR is instead provided for by the UK's Transmission Operator (National Grid) within their existing electricity balancing services portfolio. This is an issue that the National Grid recognises as causing some barriers to DSR uptake due to the participation requirements, and is trying to address this issue through their Power Response programme which aims to increase DSR usage so that it represents 30-50% of balancing capability by 2020 (National Grid, 2017). Cappers et al. (2013) note that the US market has similar issues in that DSR is provided via traditional balancing services aimed at large generation plants, which makes it difficult for uptake of DSR participation due to the entry conditions.

Market participation requirements present several key barriers for DSR. In the UK DSR has traditionally been provided from the STOR market, as reviewed in section 2.1.1. However, to participate requires meeting the entry requirements, namely through providing a minimum of 3MW, 20 minutes response time, up to 2 hours event duration, and ability to provide agreed capacity over set timeframes. These conditions can be met by the traditional large generation plants, but are more difficult for individual businesses, particularly if trying to provide turndown DSR from appliances with variable usage. Each of these conditions causes different restrictions for DSR and its uptake, for example, Grünewald & Torriti (2013) note that the 20 minute response time means that it excludes users who could take advantage of pre-loading of cooling or heating systems to cover the event period. Attempts to provide specific markets for DSR have started to emerge, with a turndown DSR only Capacity Market auction being held in March 2017 which awarded 300MW (Ofgem, 2017a). However, while the auction provided higher financial incentives (£45/kW compare to £22.50/kW for the normal auction) participants still needed to meet the standard Capacity Market rules, which require the potential for providing capacity for 4 hours or more and restricts who could viably participate.

The impact of the market's insufficient financial incentives for DSR uptake is reflected in The Energyst (2017) UK survey finding that 'Return on investment not attractive enough' is the second highest reason for non-participation. Owen et al. (2012) highlight this issue with the example of high volume manufacturing whereby any DSR interruptions cause lost production time, which translates directly into lost income. Therefore, financial incentives of DSR would need to be sufficient to offset this to enable volume manufacturers to participate.

2.2 The role of DSR aggregators

DSR aggregators are defined by the UK Office of Gas and Electricity Markets (Ofgem) as "*third party intermediaries specialising in coordinating or aggregating demand response from individual consumers to better meet industry parties' technical requirements for specific routes to market"* (Hay & Macwhinnie, 2015). Aggregators play an important role in providing DSR, with a 2017 UK survey showing 66% of participants accessing DSR via an aggregator (The Energyst, 2017). In U.S. aggregators have been deemed a key reason for the success of DSR with the example of one TSO, Pennsylvania-New Jersey-Maryland, procuring 75% of their DSR capacity from aggregators (IRGC, 2015). The aggregators' role in the UK market is highlighted in the 2015 National Grid STOR participant fuel usage report, which showed aggregators provided 82% (195MW of 237MW) of the Turndown DSR category (National Grid, 2015).

A primary reason DSR aggregators play an important role in the market is their ability to overcome a number of the barriers raised in section 2.1.4, in particulate the minimum reduction requirements for participation (SEDC, 2016). This is illustrated by the UK DSR programmes having minimum reduction requirements of 500kW to 3MW as outlined in Table 2. These minimum requirements are not readily achievable by many individual end users (expressly for Turndown DSR as noted in section 2.1.3). So, there is a need for collective participation with the support of aggregators to meet the requirements. Proffitt (2016) asserts that aggregators also help address a number of the technical and operational barriers by providing technology integration, DSR event management, and financial reconciliations required for programme participation.

As aggregators act as an intermediary between the DSR programme providers (i.e. The National Grid) and end users, their business models work on the basis of the aggregator taking a share of the value generated from the end user's participation. From the researcher's experience of working with KiWi Power, the standard market share is normally around 30% of the DSR programmes value and covers the aggregator's costs for acquiring and managing the end user while hopefully still making a profit. Aggregators are enablers for end users who want to participate but cannot meet minimum programme requirements. However, this also means that aggregators become gatekeepers for deciding if an end user can participate.

The role of aggregators in DSR is less established in literature as demonstrated by a paper search on the key terms of 'demand response', 'demand side response', and 'aggregator' in ScienceDirect and IEEE Xplore. Figure 6 shows that searching on 'demand response' or 'demand side response' in paper titles produced 1,933 papers published between 2006 to 2016, of which 149 had 'aggregator' in the title, keywords, or abstract. Figure 6 shows a recent increase in DSR research with 58% of papers being published between 2014 to 2016. The plot also shows how research on DSR aggregators is limited but also increasing. The limited research on DSR aggregators focuses on two key areas: (1) issues faced by aggregators; (2) and modelling of how aggregators fit into electrical networks.

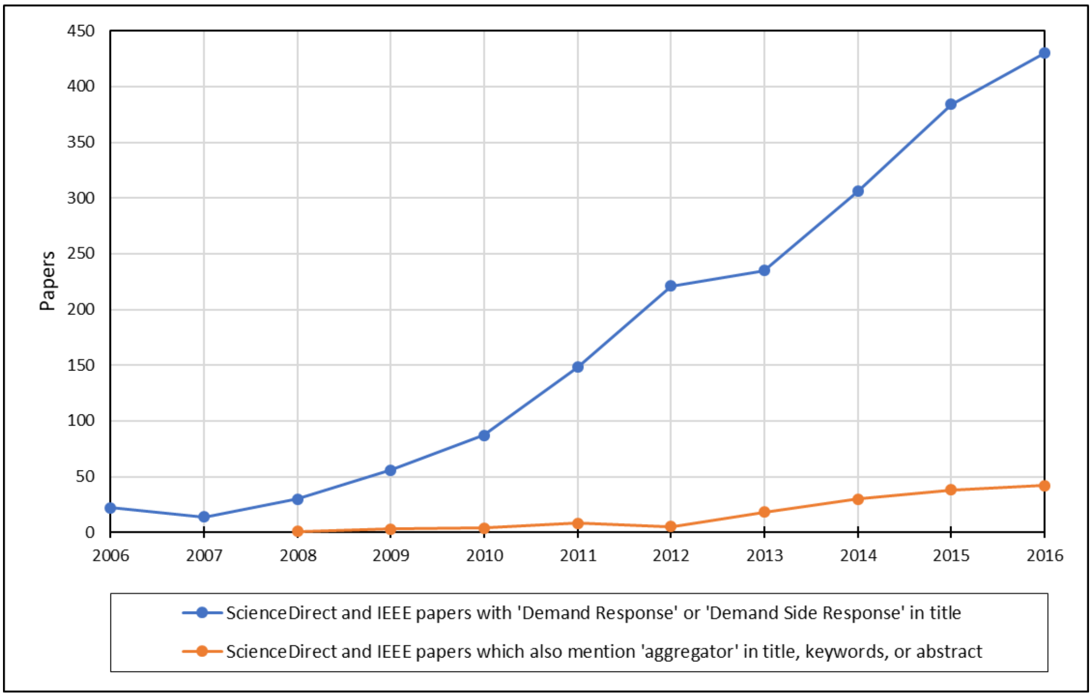

Figure 6 - Count of DSR and DSR Aggregator Papers on ScienceDirect and IEEE Xplore

2.2.1 Issues Faced by Aggregators

A common theme in the literature on aggregators is the regulatory and market access issues encountered. Research by Van Dievel et al. (2014) on regulatory frameworks and barriers of DSR in electricity distribution finds that development of DSR is restricted by current regulatory frameworks in the EU impeding the ability of aggregators to enter the distribution market. The authors provide example restrictions of fixed trading charges (such as membership and entrance fees), minimum trading volume, and minimum capacity. Lack of regulatory frameworks for aggregators is an ongoing issue with a more recent 2017 report on DSR in Europe by the SEDC highlighting this problem, along with general access to markets resulting in holding back increased usage of DSR (SEDC, 2017). In the UK, a 2015 positioning paper by Ofgem states that while there is no regular or standardised definition for aggregators, this has not prevented them entering the market. Yet its authors agree that it could hinder DSR benefits by not appropriately apportioning balancing costs when DSR is used, and that a common means of measuring performance is required.

Additional non-regulatory issues faced by aggregators are highlighted in an Ofgem report that summarised information gained about how the UK electricity system flexibility could be improved (Ofgem, 2017b). This report found that a lack of regulation and industry codes of conduct or other commonly applicable standards for aggregators has resulted in concerns in the areas of consumer confidence, transparency and reputation. Areas generating disquiet included a need for greater transparency for end users of revenues and anxiety about an apparent inability to determine the reputability of aggregators. The report also highlighted end user concerns about the aggregator systems not being robust or secure enough, with perceived risks to network stability due to cyber threats, and concerns about the potential for system errors (or risks of these) capable of causing failures to deliver capacity when required. A 'UK DSR market snapshot' by the National Grid points out that open competition is essential to maintain in the DSR market to ensure that end users have the freedom to choose a DSR supplier without being inhibited by regulations that would require them to go via their existing electricity supplier (National Grid, 2016b).

2.2.2 Aggregators' Role in Electricity Networks

The second area of literature focus is the modelling of aggregators' interaction with electricity networks. Using physically based load models and self-organising map methods Varadarajan & Swarup (2007) showed that aggregators can optimise selection of end users based on current DSR requirements to maximise benefits. By optimising end user selection for a given DSR event, instead of using a standard set of users, these authors claim that the market operation of DSR will improve, as a consequence of enabling a wider selection of users and programmes. To support residential DSR, Gkatzikis et al. (2013) propose a hierarchical market model, whereby aggregators compete with each other to provide incentives to end users, in order to modify consumption patterns. They compare this model to a standard flat-rate pricing structure, and determine that their method provides greater cost saving benefits for all parties involved. Vardakas et al. (2015) reviewed 11 DSR pricing methods and 16 optimization algorithms, and determined that the selection will differ based on the party involved and intended outcome. However, this review did not specify the optimal overall or party specific method.

2.2.3 Aggregator Knowledge Gaps

The limited literature on DSR aggregators mainly focuses on the issues that aggregators face with integrating into a traditional market and how their involvement in future smart grids. This means there is a significant gap in the literature regarding how aggregators operate, particularly when determining if an end user's operations are suitable for DSR. This is an important gap given the ability of aggregators to act as a gatekeeper for DSR take up, with this specific issue being raised in a report by the International Risk Governance Council on DSR (IRGC, 2015). The Council identified that aggregators will naturally favour large end users over marginal smaller opportunities due to the costs of operation arising where the full potential of DSR is not being realised. The first objective of this thesis aims to help address this knowledge gap in chapter 3 by reviewing how an aggregator (KiWi Power) works to understand and review the criteria used by an aggregator when determining site suitability for DSR.

2.3 DSR Site Potential Estimation

As aggregators need to ensure that applying DSR to a site will be profitable, an important aspect of the suitability assessment process is determining the site's DSR potential. Understanding a site's DSR potential requires analysis of available generator or turndown assets to try and comprehend the level of kW demand reduction possible, and for how long it can be reduced for without detrimentally impacting end users. This information is critical for two reasons. Firstly, most DSR programmes will apply penalties if contracted levels of reduction are missed. Secondly, implementing DSR has cost and time impacts. Therefore, creating a business case requires knowledge about the site's DSR potential in order to calculate the anticipated financial returns from enabling DSR participation. This section first covers why understanding a site's DSR potential is important before looking at methods used for DSR estimation.

2.3.1 The Importance of Understanding a Site's DSR Potential

The primary reason for understanding a site's DSR potential is to verify whether all requirements of the intended DSR programme can be met. As outlined in Table 2, UK DSR programmes have a number of requirements for participation that include meeting response time, event duration, and minimum electricity usage reduction targets. Failure to meet a programme's requirements could result in penalties or exclusion from participation. As an example, the UK STOR programme requires participants to provide a guaranteed kW reduction amount for set periods of time of up to 14 hours per day (National Grid, 2016c). If STOR participants underdeliver by more than 5%, then financial penalties are applied and progressively increased with the risk of ultimate removal from the programme if a participant fails in meeting guaranteed reduction levels too many times. The severity of penalties will vary by country and DSR programme. For example, the US PJM Interconnection transmission organisation has a similar penalties schemes as the STOR programme, with no penalties being applied for deviations within 5% of the scheduled amount (Burger et al., 2017). Conversely, the Spanish DSR programme is very strict, as authorities will exclude a site if it fails to meet obligations twice (SEDC, 2017). This means that correctly determining the DSR potential of a site is important for ensuring capability for continued participation in programmes in a way that maximises financial rewards and supports the profitable provision of DSR.

The second key reason for understanding a site's DSR potential is to determine if it is financially feasible to participate. Understanding a site's kW reduction potential enables financial calculations to be made based on expected returns from suitable DSR programmes. This is important given that 79% of respondents of the Energyst UK DSR survey, for example, stated that they participated in DSR in order to generate income from their assets (The Energyst, 2017). Once the potential financial returns are known, then these can be weighed against any DSR enablement costs in determining the viability of an organisation's business case for adopting DSR. The same survey also shows that the second highest reason for deciding not to provide DSR is insufficient returns (the first being unsuitable assets). This issue is highlighted by this respondent's quote about upgrading generators for a DSR programme: *"When I was talking to the aggregator my eyes lit up. But then the quote for the metering came in. For 12 sites, it was well into six figures. That was sobering. If you had that in your pocket, would you spend it on metering? It's a good opportunity but it wouldn't be the first place I would spend it."* (The Energyst, 2017). This means that failure to understand and articulate clearly the site's potential for DSR and likely financial returns arising from participating could result in incorrect or uninformed business decisions.

2.3.2 Methods for Estimating DSR Site Potential

Research into energy usage in buildings is extensive, and will be reviewed in section 2.4. However, research on application of these approaches for DSR estimation is limited. In the UK the National Grid has an end user guide for DSR, which recommends that new sites organise an onsite survey to understand their DSR potential (National Grid, 2016a). A survey will provide invaluable information about the energy used at a site, but with two limitations. Firstly, site surveys have a time and cost implication that can prevent them from occurring at all, particularly if the site is small and reliant on an aggregator to perform the survey for them. The aggregator may choose not to undertake that survey, for example, if it does not perceive enough potential to justify the cost and time involved (IRGC, 2015). Secondly, while a survey may adequately identify electricity using assets, it is not as adequate in terms of identifying asset usage patterns over time, as noted in a comprehensive survey of 2,800 commercial facilities in California (California Energy Commission, 2006). For DSR it is important to also know the asset usage over time as most programmes will require guaranteed reduction levels over set time periods that normally cover one year (for example, the UK STOR programme).

While a site survey provides important information about available assets, additional analysis is likely to be required to understand the site's DSR potential overtime. Literature on methods for this analysis is very limited, with the majority of contributions to this field originating from the Lawrence Berkeley National Laboratory - Demand Response Research Center (DRRC, 2017). Their research into DSR has covered several areas including methods for assessing the DSR potential of buildings. To help improve the assessment process, the DRRC developed the Demand Response Quick Assessment Tool (DRQAT) (Yin & Black, 2015) which uses the EnergyPlus whole building energy simulation program (U.S. Department of Energy, 2017) to predict DSR potential using predefined building models and a limited set of user selectable variables. While the DRQAT software helps to make the assessment process easier, it introduces other limitations, notably that it will only work for predefined building models, which are currently offices and retail buildings based in California. The DRRC also recognises that there are still many input uncertainties affecting the accuracy of DSR assessment, including operational behaviour and space loads, which are hard to capture in the DRQAT's model. To overcome these uncertainties, the DRQAT uses metrics of peak demand (kW), absolute demand savings (kW), and relative demand savings (%) to compensate for differences in actual and forecasted load shapes.

The DRRC have also looked into understanding the predictors that influence how well a building will perform when enabled for DSR (Mathieu et al., 2010; Piette et al., 2011). This research showed that the level of turndown potential could be linked to temperature if the DSR assets demonstrate varying levels of usage based on external weather conditions, with prediction uncertainty then being approximated by the DRRC using the standard error. However, this approach has limitations because it requires detailed asset usage information to enable correlation between weather and usage levels. Another assessment approach proposed by Panapakidis et al. (2014) is to cluster electricity usage of a building into profiles that can then be used to ascertain DSR turndown opportunities based on variance between the profiles. These authors try to reduce uncertainty by testing a range of cluster lengths to find an optimal number that minimises the overall sum of squared errors. This method has the advantage of only needing the building's overall electricity usage records. However, it is limited by the assumptions made when deciding what the differences between profiles mean in terms of individual asset usage.

2.3.3 DSR Estimation Knowledge Gaps

The review of DSR estimation methods shows there is limited research in this field. However, the extent of proprietary commercial analysis of DSR estimation methods is unknown. Information gained from one UK-based DSR aggregator, KiWi Power, implies that commercial estimation methods used by this aggregator and its competitors are less complex, and likely to have higher uncertainty levels than the ones described. The literature review highlights two key knowledge gaps: (1) that there is a lack of research on DSR estimation methods; (2) and that no comparison of the identified methods has been undertaken to explain how these differ in relation to information requirement, estimation accuracy and outcomes. The first knowledge gap will be further explored in the next section 2.4. The second objective of this thesis aims to help address the second knowledge gap in chapter 4 by performing a comparison of the outcome uncertainty of the known DSR potential estimation methods, evaluated against the level of each method's informational requirements.

2.4 Estimating Energy Usage in a Building

In contrast to the limited number of DSR estimation methods, there is an extensive range of approaches for estimating energy usage in buildings. This has been primarily driven by an increased focus by governments and industry on options for improving energy efficiency to help address climate change, as highlighted by the European Commission's 2020 Climate & Energy Package target of achieving a 20% improvement in energy efficiency by 2020 (European Commission, 2007). Lippert (2009) points out that the common management adage "You can't manage what you don't measure" applies to energy efficiency, which has resulted in methods for measuring and predicting energy usage becoming an extensively research topic. This section reviews the different energy usage estimation methods to understand how each could be used to improve DSR estimation.

Undertaking this review in a structured and comparable way requires use of a common classification framework for the many different methodologies. Before Borgstein et al. (2016) undertook a review of building energy estimation methodologies, these researchers first evaluated the classification approaches that had previously been used to create their own framework, identifying five main categories: Engineering, Simulation, Statistical, Machine Learning, and Other. This broadly follows a similar categorisation of methodologies adopted by Zhao & Magoulès (2012), except that they included 'Simulation' methods within their 'Engineering' category. Li et al. (2014) provided a contrasting approach, by using for their classification framework three main categories that each represent the level of user input required: black box (low input), grey box (medium input), and white box (high input). Due to the comprehensive review of classification approaches undertaken by Borgstein et al. (2016), and their subsequently detailed evaluation of estimation methodologies, this section uses the five categories that Borgstein et al. defined as the framework for reviewing how existing energy usage estimation methods could be utilised for DSR estimation. The review will be undertaken by describing each estimation method before then assessing its potential adoption for DSR estimation based on the ability of the method to provide the information required to determine a site's suitability for DSR as described in section 2.1.

2.4.1 Engineering Methods

Engineering energy estimation methods use simplified models comprising mathematical equations that are based on physical principles to predict usage. While there is crossover between engineering and simulation methods, Raslan & Davies (2010) explain the difference by stating that engineering methods '*generally use steady-state models that average variables over a long period in which all building parameters are fixed'* and can include '*quasi-steady states which account for the effect of some transient parameters such as weather'*. Conversely, simulation methods are described as using complex computer programmes that take all building interactions into account when attempting to simulate it. This section focuses on the simplified engineering-based methods, which fall into three general sub-categories: HVAC specific, End Use Aggregation, and Load Profiles.

HVAC specific calculation methods are deemed to be the most complex of the engineering energy estimation approaches as these have a high dependence on external factors like the weather (Borgstein et al., 2016). While different calculation methods exist, there is an International Organization for Standardization guide 52016-1:2017 (ISO, 2017). This outlines specific calculation methods, and according to Kim, Yoon, & Park (2013), these can provide very similar outcomes to using the EnergyPlus simulation tool, if calibrated correctly to the building and asset types being assessed. Other approaches include using degree-day calculations to estimate HVAC energy usage, which assumes that other energy usage in the building will be consistent throughout the year (CIBSE, 2006b). As reviewed in section 2.1.3, HVAC systems are deemed suitable for DSR due to the flexibility of these systems for being temporarily turned off without obviously impacting users. Therefore, this method offers an extensive array of HVAC specific calculation options for assessing the suitability of these systems for DSR. However, the effectiveness of this method is also restricted as it is only applicable to HVAC systems and requires a high level of information on the systems to undertake.

End Use Aggregation is a 'Bottom-up' assessment approach that provides estimations for the energy usage of all assets in a building that can then be aggregated into total building consumption. The CIBSE TM22 guide sets out a formal methodology for this technique whereby benchmarks are used to estimate each area's usage (CIBSE, 2006a). Analysis of this approach by Burman et al. (2014) showed that while it can be effective, caution needs to be taken as small deviations from the benchmarks used to estimate energy usage can compound into significant aggregated differences from actual consumption. They also note that missing benchmarks require the user to make educated guesses, and the difficulty in assessing assets with high variability like HVAC. For DSR estimation, this method offers the ability to assess all assets in a building that have low variability through usage of existing benchmarks. It would be limited by not being able to provide accurate estimations for assets with variable usage levels and risks introducing uncertainty from incorrect application of benchmarks.

Load Profiles are used to improve engineering calculations by being able to link usage of an asset to a parameter that determines which load profile or part of the profile is used to calculate the assets usage levels. This approach aims to overcome limitations of the aggregation method by allowing for variation in estimates based on a predefined parameter. Yao & Steemers (2005) demonstrated this approach by creating resident level energy usage profiles that were then applied using a sessional parameter to enable more accurate yearly predictions of energy usage. Liddiard (2014) used survey information from over 300 non-domestic buildings to produce room-scale energy usage profiles for 16 premise types. The profiles are then applied using premise type and room floor space parameters to calculate overall energy usage. Usage of load profiles for DSR would provide the following benefits once the profiles have been created: low informational requirements enabling usage at an early stage in assessing a new site's suitability for DSR, greater accuracy than more static benchmark methods due to the ability to link usage to a predetermined parameter, and ability to provide estimates for assets with variable and non-variable usage levels. Yet this method is also limited by its reliance on the availability of suitable data to create the required profiles.

2.4.2 Simulation Methods

Simulation methods utilise complex computer programmes to model energy usage in a building based on interactions between components. The usage of building energy simulation models is extensive, as highlighted by the 59 programs listed on the Building Energy Software Tools website (2017) (The list was formerly managed by the US Dept. of Energy). The ability of simulation methods to perform complex calculations enables detailed building energy models to be created for new or existing buildings. The CIBSE AM11 guide outlines usage of simulation programmes and deems them a necessity for predicting the performance of a building, referencing that these are continually improving to address increasingly complex systems and interactions (CIBSE, 2015). Yet one of the biggest issues with simulations is the performance gap between the model and actual usage. Dronkelaar et al. (2016) reviewed performance gap literature covering 62 buildings and found a deviation of 34% with a standard deviation of 55%. These authors found the primary reasons for differences were due to: specification uncertainty in modelling, occupant behaviour, and poor operational practices. Attempts have been made to overcome the gap issue through model calibration. However, a review of these methods by Coakley et al. (2014) found mixed results, and that there is no consensus or standards for which approach should be used. Building Energy Models have already been successfully used for DSR estimation as described in section 2.3.2 (Yin & Black, 2015). A key benefit of using models for DSR is their ability to provide detailed information on all energy usage in the building over extended periods of time (for example, over one year). This information can then be used to gain an understanding of the buildings DSR potential. The practical viability of using energy models for assessing DSR suitability is limited though by the extent of time and skills required to create a usable model, as demonstrated in chapter 4, and the risk of performance gaps between modelled and actual usage values.

2.4.3 Statistical Methods

Statistical methods provide 'top-down' approaches that utilise the building's overall energy usage records (e.g. monthly electricity bills, or in the UK, half-hourly meter records) to try and understand how the energy is used and to predict future usage. Zhao & Magoulès (2012) found that regression-based models were the most common statistical method used for energy prediction of buildings. As outlined in the ASHRAE Guide 14 on Measurement and Verification, regression analysis enables *'evaluation of the behaviour of the facility as it relates to one or more independent variables (e.g., weather, occupancy, production rate)'* (ASHRAE, 2015). The ASHRAE Guide recommends using regression analysis when whole building energy analysis is being performed. This enables baseline usage values to be determined that can then be used for forecasting based on the variables used (e.g. forecast next week's energy usage based on expected weather), and for measuring impacts of building changes (e.g. determining energy savings after installed a new HVAC system). As described in section 2.3.2, statistical regression models have already been successfully used for DSR estimation based on using outside temperature and prior DSR event outcomes to predict future potential (Mathieu et al., 2010; Piette et al., 2011). They are however limited by data availability which will dedicate what analysis can be performed.

2.4.4 Machine Learning Methods

Machine learning is described by Borgstein et al. (2016) as '*a term used to describe algorithms that can "learn" from data*'. The authors' review of literature about using machine learning for energy prediction identified two main sub-categories: artificial neural networks, and cluster analyses. Artificial neural networks (ANN) are a 'black box' prediction model that self-trains based on a set of data with the internal links being unknown, or specified by the user. These networks are effective for solving complex problems and can perform successfully where other models do not (Kalogirou, 2015). Neto & Fiorelli (2008) performed a comparison between ANN and building simulation models, and found the average forecasting errors were 10% and 13% respectively. Based on these results, the authors note that ANN can be an effective tool for energy forecasting if data is available for training, but the method has the disadvantage of not being able to evaluate different strategies to the extent possible with a building simulation. This means that they would be suitable for forecasting the DSR reduction ability of a site and can be continually updated with new data without requiring additional user involvement. They however are limited by available data, with forecasting of individual asset DSR reduction only being possible if specific usage data on the asset is available. The black box nature also means that users would need to trust the outcomes without being able to verify the logic.

The second category, cluster analyses is deemed an effective data mining technique for grouping similar types of objects, and is commonly used for analysis in the energy sector (Borgstein et al., 2016). Gao & Malkawi (2014) showed that clustering is able to provide a comprehensive approach to energy benchmarking. Instead of creating benchmarks by grouping based only on the traditional user-type variables, they showed that the clustering approach can be used to consider multiple factors that enable a better grouping of buildings based on similarity of energy usage. Care however needs to be taken when using clustering as Hsu (2015) shows the results can vary dramatically depending on '*choice of algorithm, variables selected, initial assumptions, and the natural shape of the data*'. As described in section 2.3.2, clustering has been utilised by Panapakidis et al. (2014) to demonstrate that it can be used to create profiles that could potentially be used for DSR estimation. It is limited to identifying large energy using assets due to requiring large differentials in electricity usage and, depending on the building energy usage, it may not be able to differentiate individual assets.

2.4.5 Other Methods

Reviews of estimation methods identify a number of methods that do not clearly fit into one of the previously reviewed categories (Borgstein et al., 2016; Zhao & Magoulès, 2012). These include disaggregation, auditing, and sub-meter methods which are covered in this section. The first other method, disaggregation, is described by Carrie et al. (2013) describes as the '*extraction of appliance level data from an aggregate, or whole building, energy signal, using statistical approaches'*. Disaggregation's primary benefit is the lower cost and time requirements of only having to monitor one single feed (for example, the main incoming electricity line or floor level feeds) instead of installing individual monitors for each appliance. However, usage of disaggregation has been limited to domestic buildings as the pattern recognition algorithms have difficulty when there are multiple appliances of the same category (e.g. computers in an office) (Rodriguez, Smith, Kiff, & Potter, 2016; Schmidt & Bansal, 2017). For DSR the potential of disaggregation enabling cheaper and easier to install monitoring methods would provide individual asset usage data to improve DSR estimation and on-going monitoring. It is currently limited by only being available for residential buildings as commercial building level solutions are still in the research stage.

The second other method, energy auditing, is described by the CIBSE Guide F (2012) as '*investigations of site energy use aimed at identifying measures for cost savings, energy savings and reductions in environmental emissions'*. Energy audits are primarily focused on finding savings (CIBSE notes a potential energy cost saving of between 10%-60%), but also provide the benefits of identifying how energy is used, which can then be used for other purposes like quantifying DSR potential. Yet energy audits are limited as these only provide a snapshot of current energy usage. This can be useful when identifying savings and then comparing energy changes after efficiency works have been completed. However, audits are less useful for forecasting purposes given the need to use additional analysis methods to manage asset variability. Energy audits are also expensive to undertake, with the CIBSE Guide F stating an expected cost of between 3% to 5% of the sites yearly energy cost to perform. For DSR an energy audit could provide the detailed understanding information of electricity usage in a building that can be used for detailed DSR estimation and to confirm enablement costs. It would however be costly to undertake and only provides a snapshot of usage levels, which can be problematic for estimating the DSR ability of assets that vary with usage across the year.

The third other method, sub-metering, involves the metering of individual assets or areas in a building. A Better Building Partnership (2011) report on metering notes how having detailed usage records from sub-meters (as opposed to the sites overall usage meter) enables analysis to be undertaken to identify efficiency improvements, alerts of incorrectly operating assets, and validates energy billing. Jones (2012) however notes a number of issues with sub-meters, including: incorrect placement resulting in absence of clarity about what is being metered, too much data that confuses analysis, and metering not being connected properly or serviced, resulting in the potential for inaccurate data. For DSR sub-metered data provides a very clear profile of an asset's usage over a year that can be used to determine DSR suitability. Usage is however limited by a lack of installed sub-meters and difficulty in accessing data, which means that sub-metering cannot be relied on as being a readerly available source for DSR estimations.

2.4.6 Building Energy Estimation Method Knowledge Gaps

This section has reviewed ten energy assessment methods of which four (Simulation, Statistical, Clustering, Energy Audit) have previously been applied in assessing the DSR potential of buildings. Six potential methods (HVAC Calculation, End Use Aggregation, Load Profiles, ANN, Disaggregation, Sub-Meter) have no known existing research on their applicability for understanding the DSR potential of a building. While each of these six methods could be used for DSR estimation, the review showed that some lend themselves to being more suitable. To understand their level of potential suitability for DSR estimation this research has classified each of six methods into three groups based on the review outcomes – least suitable, might be suitable, highly suitable.

- **Least suitable** are the disaggregation and sub-meter methods. Both methods provide detailed usage information about individual assets which would enable DSR estimation. However, each method is deemed 'least suitable' due to lack of actual uptake. Disaggregation technology is currently limited to residential buildings, with non-domestic applications still in the research phase. Sub-metering is limited given the lack of installed meters and difficulty in accessing data from ones that have been installed. (Jones, 2012).

- **Might be suitable** are the HVAC Calculation and ANN methods. The specific focus of the HVAC calculation method limits its suitability for DSR estimation of a building, as other methods will also need to be required to gain a complete profile. ANN methods are promising in their ability to accurately forecast energy usage. However, ANN are limited as these networks require substantial amounts of data for training, focus on providing forecasts of anticipated site (and not asset) level energy usage, and their 'black box' nature can result in difficulty with understanding the logic used to generate the outcomes.

- **Highly suitable** are the End Use Aggregation and Load Profile methods. The End Use Aggregation method uses existing benchmark information about assets to provide an overall indication of a building's DSR estimation potential. This method is more suited to fixed usage assets as it can have difficulty with estimating variable usage assets. The Load

Profile method helps overcome the End Use Aggregation method's inability to estimate assets with variable usage levels by linking a suitable parameter to the performance of the variable asset. To create the link requires suitable data that enables an asset's usage levels to be linked to a parameter. Access to this data (or lack of access) can limit the ability to use this method.

While suitability judgements are subjective, this research has grouped estimation methods based on the reviewed research and each method's perceived ability to provide DSR estimations. With knowledge gaps existing on usage of these six methods for DSR estimation, this suitability grouping enables targeted focus on which ones should be researched first. Ideally, all six methods would be assessed for how each could be used for DSR estimation. However, prioritisation is required given limited research time, and the importance of ensuring that this thesis provides an adequately thorough, in-depth research study. Therefore, the third and final objective of this thesis focuses on using the Load Profile method in chapter 5 to develop and evaluating a demand response estimation method which is capable of using asset usage profiles to reduce the uncertainty and cost of DSR estimation.

2.5 Chapter Conclusion

This chapter has reviewed existing literature on DSR to identify the knowledge gaps that will be addressed by the research objectives of this thesis. A broad review of DSR and its usage showed that focus on DSR is increasing, primarily due to the need for government and industry to meet the future needs of a low carbon economy. Yet reviewing research into the challenges of providing DSR identified the many barriers that are limiting the expansion of DSR, validating the need for continued research into DSR and opportunities for improving its adoption. The review then focused specifically on literature applicable to each of the research objectives.

For the first research objective, the review focused on literature about DSR aggregators. This review found very limited existing research on aggregators, with the existing literature mainly focusing on issues faced by aggregators and modelling how aggregators fit into electrical networks. There was almost no research found covering how aggregators actually work or how they decide whether a site is suitable. This outcome confirms that the first objective '*To map out the criteria used by an aggregator to determine site suitability for DSR'* will address an important gap in existing knowledge on DSR aggregators.

For the second research objective, the review aimed to understand the publicly available literature on existing DSR estimation methods. This identified three methods. No existing comparisons for these methods were found either, which prevents the ability of users to select a method based on knowing how they compare in terms of input requirements and expected levels of output error. Therefore, the second objective '*To perform a comparison of the outcome uncertainty in DSR potential estimation methods, evaluated against the level of informational requirements of those methods'* will address this knowledge gap.

The final research objective sets out to address the lack of DSR estimation methods by developing a new approach, which required an assessment of potentially available energy usage estimation methods. The review of existing building energy usage methods identified six major categories of methods which have been directly used for DSR estimation. The third objective '*To develop and evaluate a model that uses asset usage profiles to reduce the uncertainty of DSR potential estimation during an aggregators assessment process'* will adopt and then adapt the Load Profile method explained in this literature review to develop a new DSR estimation approach.

3 DSR Aggregators: How They Decide Customer Suitability

This chapter addresses the first research objective *'To map out the criteria used by an aggregator to determine site suitability for DSR'*. As outlined in section 2.2, the existing research on DSR aggregators is limited. Yet aggregators play a key intermediary role in facilitating uptake of DSR by combining small flexible loads from multiple end users into a virtual single load that can then be managed by the aggregator in meeting the criteria of system operator programmes. If aggregators are to deliver a significant upscale in DSR penetration, it is important to understand how they determine the suitability of an end user for DSR, and the challenges that end users may face during the suitability assessment process. Currently there is no known research on aggregators' end user suitability assessment. The limited research that does exist is focused on the analysis of the technical and economic models of their integration into electricity markets (Gkatzikis et al., 2013). This chapter addresses this gap by increasing the body of knowledge about the aggregator's role in the adoption of DSR through examination of KiWi Power's new client assessment process.

This research involved access to only KiWi Power and its data for determining how DSR aggregators assess new sites. This single source of data for new site assessments limits the ability to directly compare the outcomes from KiWi Power's processes with those of other aggregator to ascertain if KiWi Power are representative of DSR aggregators in the UK and internationally. However, companies operating in the DSR aggregator market can be classified into comparative groups, which provides context for identifying how KiWi Power compares and whether other companies are likely to have equivalent site assessment processes. Figure 7 provides a phase space schematic of DSR aggregators that includes KiWi Power, illustrating how KiWi Power can be classified in the context of other aggregators. The x-axis indicates an aggregator's sales approach. An 'In-Direct' sales approach means that sites are acquired through partnerships (for example, embedding the aggregator's DSR control into a device that is then sold by the manufacturer with no guarantee that the customer will utilise the device's DSR option). A 'Direct' sales approach means sites are directly contacted by aggregators to ascertain if the site owner wants to undertake DSR (for example, an aggregator's sales team calls the owner/manager of a building to try and sign them up for participating in a DSR programme). The y-axis indicates the sizes of assets that the aggregators target. Plotting aggregators using these axes identifies in Figure 7 a correlation between the sales approach and asset size. Direct approaches aim to obtain large assets to reduce a site's per MW acquisition costs, whereas the in-direct sales approach aims to have a low per asset acquisition cost, which enables smaller sized assets to be used to form an overall DSR portfolio. In this analysis, KiWi Power emerges with others as an aggregator that adopts a direct sales approach for larger assets.

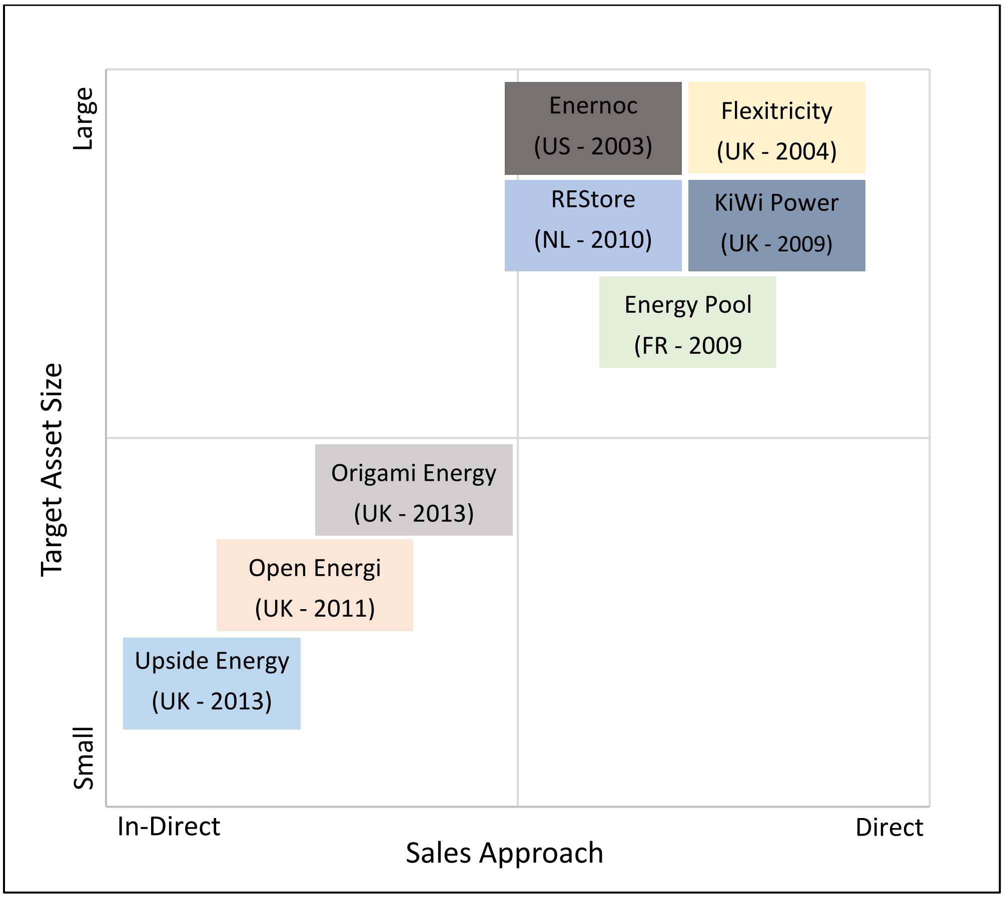

Figure 7 - Comparison of DSR Aggregators in Terms of Sales Approach and Targeted Asset Size (With the brackets showing head office location and establishment year)

The selection of aggregators and their positions on Figure 7 is speculative but informed by an understanding of the DSR industry and its trends, reflecting that organisations that entered the market early have followed the traditional approach of direct acquisitions, whereas newer organisations target gaps in the market through focusing on smaller asset sizes. KiWi Power sits in the group of traditional DSR aggregators that focus on using the direct sales approach. Therefore, the results from this chapter's examination of KiWi Power's new client assessment process are relevant for traditional aggregators that pursue direct sales and large assets, which is how a large proportion of the DSR market currently operates in the UK, as well as France, the Netherlands and North America.

In this chapter, Section 3.1 describes the analysis approaches used to understand KiWi Powers' new client assessment process. Section 3.2 sets out the analysis results and discusses these findings. Section 3.3 concludes this chapter's findings.

3.1 Analysis of KiWi Power's Client Assessment Process

To understand the DSR decision-making process and what comprises a 'suitable client' (also referred to as 'end user' or 'site') requires a systematic review of the KiWi Power's new client assessment process. To undertake this review, two streams of investigation were undertaken. Firstly, a detailed analysis of KiWi Power's new client assessment workflow system, in order to gain an in-depth understanding of their client base and the reasons why clients are or are not categorised as suitable for DSR. Secondly, based on the workflow system review, semi-structured interviews with KiWi Power employees, in order to gain further insights into each stage of the workflow process and the reasons behind client's suitability assessment outcomes. Section 3.1.1 provides an overview of the KiWi Power new client assessment process as background before section 3.1.2 reviews the methods used to analyse the workflow system and section 3.1.3 reviews the methods used to interview the employees.

3.1.1 Overview of New Client Assessment Process

KiWi Power works on a shared revenue model, whereby this aggregator takes a percentage of the DSR revenues earned by the participating client's site in return for managing their DSR service. Their share of the revenue covers many costs, including: normal business operating costs (office rent, utilities, support staff), hardware and contractor costs to enable the client's site for DSR, acquisitions costs of the sales and engineering teams, and investment costs in research and development. Therefore, the aggregator requires a robust assessment process for new clients, to ensure their site will be suitably profitable before being accepted for DSR if the aggregator is to remain solvent. The new client acquisition process undertaken by KiWi Power is split into three key stages, and managed through a Customer Relationship Management (CRM) workflow program called Salesforce (Salesforce, 2016). Each stage is reviewed in detail and informed by KiWi Power training guides and informal discussions with employees to understand the process.

Stage-1 of the acquisition process is called *'Initial Client Contact'* and normally takes less than a month to complete. This stage starts once a sale team member has made an initial contact with a potential client or lead, and has undertaken a high-level assessment to determine if there is any site potential. This assessment is normally performed via a phone call with the client's facility manager, or a similar person who has knowledge about the site's power usage and operations. During the phone call, the salesperson explains the concept of DSR and what KiWi Power does before asking specific questions on two areas: backup generators and large electricity assets (see sections 2.1.2 and 2.1.3 for background information on suitable DSR assets). The sales lead starts by asking the client if they have a backup generator. If yes, then further questions are asked, including: what is its kW capacity, how old is it, is it tested regularly, and do they have a G59 connection (see section 2.1.2 for more information on G59 connections). The client contact is then asked about electrical assets that could be used for Turndown DSR. If the client has any suitable assets, then the aggregator will try to obtain from the conversation an initial idea from the client about the types and quantity of assets, kW rating sizes, and how they are currently used, in order to get a sense of the potential for temporarily turning the identified assets down or off. The sales person then decides, based on their training and prior experience, if this client is worth pursuing further. If yes, then the aggregator progresses to stage 2.

Stage-2 of the acquisition process is called *'Site Assessment',* and normally takes between one to three months to complete. It primarily covers qualification of the client's site for DSR and contractual agreements. At this stage, the client's site undergoes a detailed assessment, starting with obtaining the site's half-hourly electricity usage records for the past year and a detailed list of all DSR-applicable generator and turndown assets. The half-hourly electricity usage records are analysed using the KiWi Power Analysis Tool to understand the site's baseload usage levels for different DSR programmes. The site's DSR-applicable generator and turndown assets are then compared against the baseload to ascertain their potential utilisation to determine the overall DSR potential (for example, if the site has a 2MW generator but only a 1MW baseload with no export ability then the DSR potential is limited to 1MW). If generators are potentially in-scope, then the aggregator checks to see if the client site already has a G59 connection and whether there is an existing export agreement in place. If not, then checks are performed to understand if the local grid can support additional export, and quotes are sought on enabling the generator to export. If the sales team believes the site has potential, then a technical review for suitability is undertaken by a KiWi Power technician via a phone call or site visit. The aggregator's technician will aim to verify that the site's assets can be used for DSR based on a number of factors, which can include: how the assets are controlled, impacts of the assets on business operations, access for installing DSR controls, mobile data signal strength, and the age and reliably of the assets. If deemed technically suitable, the last step of this stage is to finalise financial agreements and undertake contract negotiations. Once an agreement is signed, the client engagement progresses to stage 3.

Stage-3 of the acquisition process is called *'Site Enablement and Go Live'* and can take between three to six months to complete. This stage shifts the responsibility of the client from the sales team to the aggregator's technicians, who are responsible for enabling the site for DSR. A key task is the installation of KiWi Power's custom-made control and monitoring equipment, which is called the Power Information Pod (PiP). The PiP provides the necessary link between the client's assets and KiWi Power's control systems. It normally gets connected to the site's Build Management System (BMS), or directly to the asset. This provides basic DSR event signals to the BMS or asset in the form of an 'on' or 'off' command. Its normal state is off, but when an event occurs, a signal is sent to the PiP that then activates the interface to send the on command. The client's assets which are linked to the PiP will then respond as required to provide demand reduction (either via asset turndown, or generator turn on). The PiP also has the capability to monitor power usage and acts as a sub-meter. As a consequence, when a DSR event occurs, KiWi Power is able to monitor the response and report back to the DSR programme initiator the amount of demand reduction achieved. Installing the PiP can be one of the challenges that extends a site's DSR enablement time, as it normally requires using the client's own electrical contractors, having access to the assets (which can sometimes require turning them off), and running new cables for control, power and signal lines. If a generator is being enabled for DSR then additional time is normally required if the G59 connection is not already in place, or where the site is still waiting for exporting approval. Once all equipment is installed, the site undergoes spot testing. The spot test will be undertaken by KiWi Power without the site knowing beforehand to verify that they are operationally fit to go live. When the site has passed all tests, KiWi Power then assigns the client to an appropriate DSR programme, and the site is officially then deemed to go live, at which point the acquisition process is complete and responsibility for the client is handed over to the operations team.

3.1.2 Methodology Used to Analysis the Workflow System

The first investigation stream performs a quantitative analysis of the Salesforce system that manages the KiWi Power client assessment process. The purpose of this analysis is to gain a detailed understanding of how clients progress through all three stages. The primary outcome of this analysis is to understand how many clients are lost at each stage, and the reasons why. This information forms the quantitative bases of this chapter's research, and informs the creation of semi-structured interviews questions for the next analysis phase, described in section 3.1.3.

The content analysis method was selected to undertake the review of Salesforce records. This method provides a proven approach to undertaking a quantitative systematic review of content to draw out meaningful insights. Bryman (2016) outlines the advantages of this research method as: it is transparent and replicable, it enables research of a timespan of records, it is adequately flexible to fit the type of content being reviewed. The disadvantages though are: the analysis is only as reliable as the records being reviewed (*research mitigation* - access to the Salesforce administrator and sales team enables clarification of missing information), the coding methods can be subjective to the coder (*research mitigation* – coding schemes were checked by fellow researchers to ensure consistency), it can be hard to gain in-depth understand of specific events in the material without talking to the creator (*research mitigation* – the semi-structured interview approach outlined in the next section 3.1.3 was specifically chosen to overcome this limitation and gain additional insights).

The data used for this content analysis was a static extract of all client records from the Salesforce database. The static extract was taken in July 2016 and contained 772 client records from 2013 onwards, although the majority were recorded in 2015 and 2016. The coding schedule was formed of three fields, as outlined in Table 4. Using this coding schedule, each client in the database was first assessed to determine its current workflow stage and whether the site opportunity had been lost, was still progressing, or was now live. Sites were deemed lost if they were marked in the system as no longer being progressed and closed. The records for each lost site were then analysed to determine a reason for why the site did not proceed any further. Once all client records had been assessed then the following analysis steps were undertaken, and the results recorded in section 3.2:

1. **Step 1 – Calculate Stage Level Percentages:**
    a. For each stage, the number of clients per status value, as defined in Table 4, was counted.
    b. The counted values were converted into two percentages: an overall percentage based on how many of the 772 clients were at each stage's status, and a stage percentage that shows the outcome of that specific stage.
    c. The percentages were then used to create the Sankey flow diagram in Figure 8.

Table 4 - Content Analysis Coding Schedule

Field	Coding Options	Coding Source and Description
Client ID	Unique identifying ID of the client record	Based on the client's Salesforce record
Status	Selection Options: 1. Stage 1 - Lost 2. Stage 2 - Lost 3. Stage 2 - In Progress 4. Stage 3 - Lost 5. Stage 3 - In Progress 6. Stage 3 - Live	Each client's Salesforce record has a current stage and status field. This information is coded by using the stage number and then determining the status. The status field is reviewed to determine if the client has been lost, still progressing through the current stage, or has gone live.
If lost, what is the reason?	The selection options were updated as new reasons for being lost were found during the analysis. Table 5 shows the final list of reasons sites were lost per stage.	The reasons for sites being lost were derived from the Salesforce status field for each client, as defined in Table 5.

2. **Step 2 – Calculate the 'Reasons for Being Lost' Percentages per Stage:**
    a. The number of occurrences for each reason a site was lost from Table 5 were counted.
    b. The overall assessment process percentage occurrence was calculated for each reason a site was lost, and included in the ranked Table 7.
    c. The reasons were then split by stage and the percentage of reasons per stage were calculated and used for comparison with the interview outcomes from the following section 3.1.3.

Table 5 - List of Client Reasons Sites Were Lost by Stage

Stage	Reasons Sites Were Lost	Coding Reason Based on Salesforce Status Field
1	Not Interested in DSR	Client showed no interest in DSR after being contacted.
1	Not Interested due to impact concerns	Client expresses concerns about impacts to business after being contacted about DSR.
1	Not Interested due to insufficient financial returns	Client not interested in DSR due to insufficient financial returns after being contacted.
1	Assets deemed unsuitable for DSR	After talking with client, the salesperson does not believe there are any suitable assets that could be used for DSR.
1	Already with another DSR provider	Client was contacted but already with another DSR provider, and not interested in changing.
1	No assets large enough for DSR	Client was contacted and after reviewing DSR assets none were deemed to have a high enough kW potential.
2	Lost interest in pursuing DSR	Client loses interest during qualification and contract steps for an unknown reason.
2	Lost interest due to internal priorities	Client loses interest due to other business priorities being deemed more important.
2	Lost interest due to impact concerns	Once the additional analysis has been performed the client believes the business impact is too high.
2	Went with another DSR provider	Client was assessing DSR provider options and decided to go with another provider.
2	Upgrade costs to enable DSR too high	After performing the detailed assessment, it was determined that the enablement costs were too high to justify continuing.

2	Assets deemed technically unsuitable for DSR	After performing the detailed assessment, the DSR assets were deemed unsuitable due to technical reasons.
2	Lost interest due to insufficient financial returns	After performing the detailed assessment and finalising financial returns the client decides they are not sufficient to justify continuing.
2	DSR potential too low for aggregator to pursue	After performing the detailed assessment KiWi Power deems the kW potential is too low to justify continuing.
2	Missed deadlines for specific DSR programme	The client was aiming for a specific DSR programme and was unable to become enabled in time to participate.
3	Assets deemed technically unsuitable for DSR	During installation or testing the assets were deemed unsuitable for technical reasons.
3	Client decided to stop progressing with DSR	During installation or testing the client decides to not continue with DSR for unspecified reasons.
3	BMS too old for integration	BMS used to control the DSR assets is found to be too old for interfacing with the PiP.
3	Turndown reduction too small once tested	Once enabled the testing identifies that the actual turndown ability is smaller or more variable than assessed, and therefore deemed unsuitable.
3	Missed deadlines for specific DSR programme	The client was aiming for a specific DSR programme and was unable to become enabled in time to participate.
3	Enablement costs deemed too high	Once enablement started unexpected costs identified that make it unsuitable to continue.
3	No generator export ability and site load to small	Site load too small or variable and without export agreement, making it unsuitable for DSR.

3.1.3 Methodology Used to Interview the KiWi Power Employees

The second stream of investigation was undertaken through a semi-structured interview of 12 KiWi Power employees. The employees covered all three stages of the workflow process, with 9 employees being focused on sales - stage 1 (initial client site contact) and stage 2 (site assessment), and 3 employees being focused on technical site enablement, and making the site operational - stage 3 (site enablement and go live). Following the University of Reading's ethics guidelines, approval was obtained to undertake the interviews with all electronic data being held on a password protected University owned laptop and all physical consent forms and paperwork being stored in a locked cabinet. All responses were anonymised with each employee being assigned a 3-character code based on their role and seniority, as per Table 6. The lower numbers of stage 3 technical employees arise because most onsite physical installation is undertaken for KiWi Power by third-party contractors with the KiWi Power employees purely focusing on managing the process.

Table 6 - Interviewee Code's and Roles

Interviewee Code	Role and Seniority
SJ1	Sales - Junior
SJ2	Sales - Junior
SJ3	Sales - Junior
SJ4	Sales - Junior
SI1	Sales - Intermediate
SI2	Sales - Intermediate
SI3	Sales - Intermediate
SS1	Sales - Senior
SS2	Sales - Senior
TJ1	Technical - Junior
TS1	Technical - Senior
TS2	Technical - Senior

The semi-structured interview was designed around each of the 3 workflow stages (see Appendix A for the interview guide, and Appendix B for the reference sheets used). For each stage, the interview was split into 3 sections. The first section consisted of providing the interviewee with a printed list of each stage's workflow tasks with a 5-point Likert scale. The interviewees were then instructed to use the scale to indicate how important the interviewee perceived each task when deciding the suitability of a site for DSR, using a scale where 1 signified 'unimportant' through to 5 signifying 'very important'. The second section provided the interviewee with a printed list of reasons that sites were lost during that stage (as per Table 5), based on the workflow system analysis in the previous section. Each reason was provided with a 5-point Likert scale with the interviewee being asked to indicate, based on their experience, the likelihood of the reason causing a site to be lost, with 1 being 'very unlikely' through to 5 being 'very likely'. The final section was a series of open questions about the workflow stage, which aimed to gain insights into an interviewee's reasons for selecting specific scale ratings and what the interviewee considered important when deciding on the suitability of a site during that workflow stage. This also included specific questioning about what the interviewee considered to be a minimum DSR kW potential for a site to be worth progressing for turndown or generator-based DSR.

The semi-structured interview method was selected over the structured and non-structured alternatives due to the benefits it presented for this research. The key benefit being the ability to incorporate mixed methods, covering both quantitative and qualitative approaches (Bryman, 2016). This allowed for the Likert scale to be incorporated to obtain quantitative data regarding the importance of each workflow task, which could then be expanded on using qualitative questioning to understand the reasoning behind the scores giving by the interviewee. The qualitative question element also provided the benefit of flexibility for asking additional questions about each stage that were not covered by the process task list (like questioning on minimum kW levels for DSR). The final advantage of using this method was the ability to manage interview timing through predefining the interview guide. As the interviews were undertaken during business hours with time-pressured interviewees, having a well-defined interview guide ensured the interviewees remained on topic and enabled interviews to be completed on time without being cut short. This method does have disadvantages though, including having to manage '*demand characteristics'* whereby interviewees will adjust their responses to meet a preconceived idea about the research outcomes (Nichols & Maner, 2008). As this interview was undertaken in a workplace setting that discussed the interviewee's job tasks, there was a high probability that answers would be skewed towards positive responses where, for example, interviewees thought that management might read the outcomes and associate these with an interviewee. Research by Nichols & Maner (2008) on this issue found that addressing it is difficult, and recommended that the interviewer remain as neutral as possible, that confidentiality of the results is reinforced to the interviewee including in any consent process, and that the results analysis should attempt to identify any patterns that may have been caused by this. Additional disadvantages noted by Bryman (2016) of the semi-structured interview method include it being a time-consuming approach (*mitigation* – only 12 interviewees were available), and being more difficult in comparison to structured interviews to compare interviewee results (*mitigation* – the usage of the Likert scale enables direct comparison, with the qualitative questioning used to support reasons behind the Likert outcomes).

The semi-structured interview included use of the structured Likert scale to understand the importance of workflow tasks in determining site suitability. The Likert scale has many benefits as a quantitative survey method that includes: (1) it is an established method that is recognisable and easy to understand by interviewees, (2) the coding scale enables detailed statistical analysis of the results, (3) its non-binary scale provides interviewees with the freedom to express varying levels of response, and (4) the scale comprises a time efficient method of obtaining multiple interviewees responses on a subject matter (Bryman, 2016; Hartley, 2014; Hasson & Arnetz, 2005). The Likert scale also has limitations that need to be addressed. The statistical analysis of Likert results was reviewed by Harpe (2015) who recommends that the median and quartile measures should be used over mean and standard deviation measures as the latter will not accurately represent values that are interval data. The number of points to use on the scale is open to debate, with Willits, Theodori, & Luloff (2016) looking at the research on this topic and concluding that while different scales could be used to gain finer levels of response (notably 7 and 9 point scales), the standard 5 point scale is deemed suitable given no clear benefits from using larger scales. Finally, the Likert scale is open to interviewee bias that can result in either an end-aversion bias whereby the interviewees will only select values from the middle options (Hasson & Arnetz, 2005), or a social desirability bias whereby the interviewees select more positive results due to either wanting to appease the interviewer, or based on their preconceptions of questions being asked (Furnham, 1986).

Once all 12 interviews had been completed the quantitative results were analysed using the following steps, and the results recorded in section 3.2:

1. **Step 1 – Likert Scale Conversion:**
    a. All the interviewee's Likert scale results for each stage's workflow tasks and reason for sites being lost were recorded into a spreadsheet (see Appendix C).

    b. Based on Harpe's (2015) recommendation, the median and quartile measures were calculated for each workflow task and reason for site loss.

    c. The median and quartile results were plotted on a 5-point scale for each workflow task and reason for site loss.

    d. All workflow tasks were ranked based on the median values and the results recorded in Table 8.

    e. The workflow tasks were also separated by stage into Table 9, Table 11 and Table 13.

    f. The reason for site loss percentages from the workflow analysis section 3.1.2 were combined with each stage's Likert scales and separated by stage into Table 10, Table 12 and Table 14.

2. **Step 2 – Boxplot Creation for Interviewee Minimum DSR kW Potential Requirements for New Sites:**
    a. Each interviewee's minimum DSR kW potential requirements for Individual Asset Turndown, Overall Site Turndown, and Generator categories were recorded into a spreadsheet, as shown in Appendix C.

    b. The minimum, lower quartile, median, upper quartile, and maximum values for each category were calculated.

    c. The calculated values were used to create a boxplot, as per Figure 9.

3.2 Results of New Client Assessment Process Analysis

The results of the assessment process analysis are reviewed and discussed over six sections. Section 3.2.1 reviews the status outcomes at each stage of the workflow, and includes a ranked list of reasons that sites were lost. Section 3.2.2 reviews the Likert scale importance rating giving to each stage's tasks by the interviewees. Sections 3.2.3 to 3.2.5 cover the three workflow stages in depth, in order to gain insights behind the importance rankings giving by interviewees. Finally, the minimum DSR kW requirements given by interviewees for sites to be found suitable are reviewed in section 3.2.6.

3.2.1 Review of New Client Assessment Process Outcomes

The Sankey flow diagram in Figure 8 shows the Salesforce database analysis results as percentage outcomes for each individual stage based on a snapshot of 772 client records that encompassed information recorded between 2013 and July 2016. For stages 2 and 3, the additional percentage value in brackets is provided to show the overall impact of that outcome. As the There are no 'still being assessed' clients during stage 1 (Initial Client Contact), as new clients are only entered into the Salesforce system once a decision has been made on if they will progress to stage 2 or are deemed lost. The results show that overall only 14% of the clients have completed the workflow process and gone live, with 33% still being progressed and 53% lost. Stage 2 has significantly higher client losses at 35%, followed by stage 1 and then stage 2 at 19% and 18% respectively. While the number of live clients seems low, if we assume the number of clients still being progressed will complete with the same level of success rate documented at stages 2 and 3, then the overall live percentage would increase to 34%.

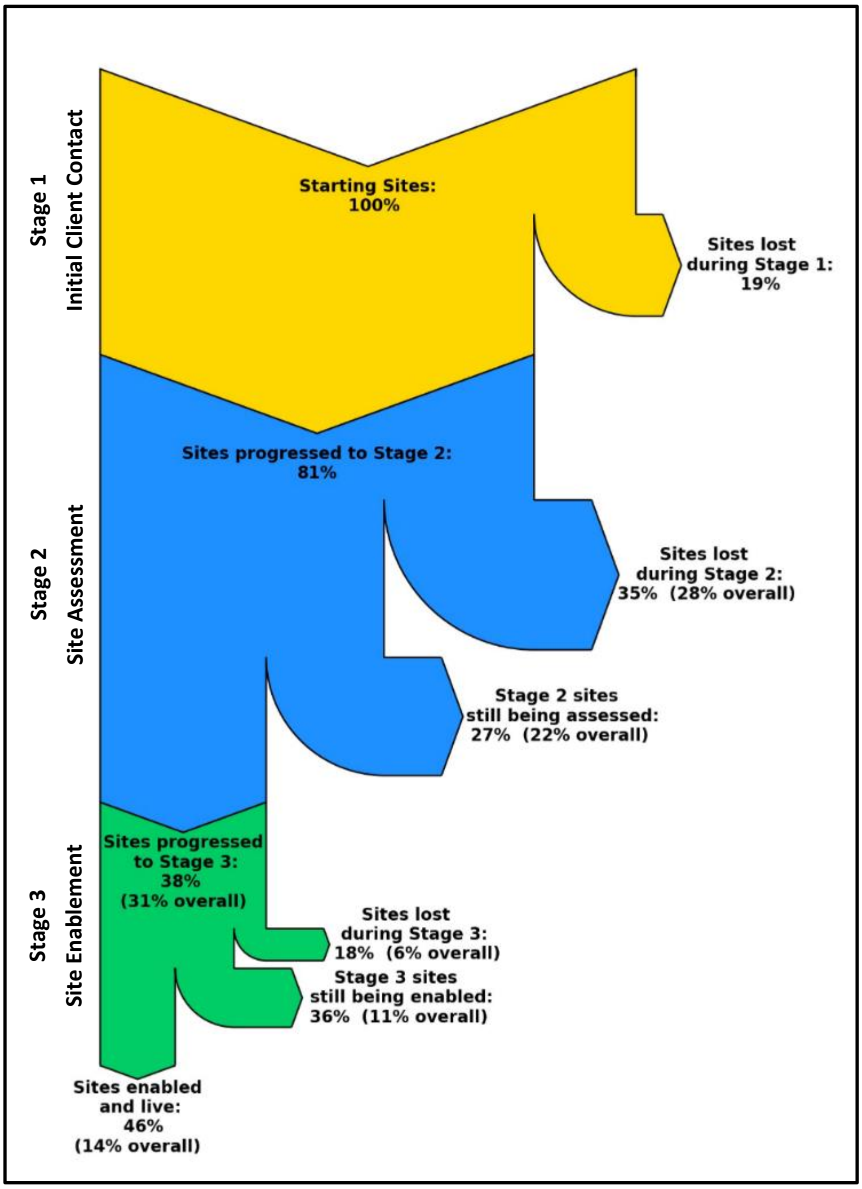

Figure 8 - Sankey Flow Diagram Showing Outcomes at Each Stage of the Acquisition Process

A ranked by percentage list of reasons why the remaining 66% of clients are lost is shown in Table 7. As only one lost reason was recorded in the KiWi Power Salesforce database, a client could potentially have also been lost for another reason but not recorded. Therefore, the outcomes shown at Table 7 are not mutually exclusive. A client might have dropped out as a participant due to being 'Not interested in DSR', for example, and not reached the next stage of assessment for asset suitability that might otherwise have also identified that the client actually had 'No assets large enough for DSR'. The table shows that the top two reasons 'Lost interest in pursuing DSR' and 'Not interested in DSR' account for 35% of losses where clients decide not to pursue DSR. The 'not being interested' reason was not identified in the section 2.1.4 review of existing DSR barriers to participation. The identified barriers in the literature that might relate to this reason are 'impact concerns' and 'not being a priority'. However, as both have been captured in Table 7, they are unlikely to be the cause which means that 'not interested' is a new barrier that is missing from existing research findings. The split between which party decides to stop pursuing DSR shows that 72% of the time it is the client's decision. The results show that the general theme of losing interest accounts for 52% of clients deciding not to continue, followed by 14% being lost due to competition. Conversely, the reasons that KiWi Power decides not to pursue a client are primarily due to the DSR assets not being suitable due to size, upgrade costs, or other technical reasons. Further insights into these reasons for sites being lost are provided during the detailed review of each stage in sections 3.2.3 to 3.2.5.

Table 7 - Ranked List of Reasons for Site Lost

Stage	Reason for Site Being Lost	Aggregator or Client Decision	Recorded Frequency
2	Lost interest in pursuing DSR	Client	21%
1	Not interested in DSR	Client	14%
1	No assets large enough for DSR	Aggregator	7%
1	Already with another DSR provider	Client	7%
2	Went with another DSR provider	Client	7%
2	Lost interest due to internal priorities	Client	5%
1	Assets deemed unsuitable for DSR	Aggregator	5%
2	Lost interest due to impact concerns	Client	4%
2	Upgrade costs to enable DSR to high	Aggregator	4%
2	Assets deemed technically unsuitable for DSR	Aggregator	4%
3	Assets deemed technically unsuitable for DSR	Aggregator	3%
2	Lost interest due to insufficient financial returns	Client	3%
1	Not Interested due to impact concerns	Client	2%
2	DSR potential too low for aggregator to pursue	Aggregator	2%
1	Not Interested due to insufficient financial returns	Client	2%
2	Missed deadlines for specific DSR programme	Client	2%
3	Client decided to stop progressing with DSR	Client	2%
3	Building Management System too old for integration	Client	1%
3	Turndown reduction to small once tested	Aggregator	1%
3	Missed deadlines for specific DSR programme	Client	1%
3	Enablement costs deemed too high	Aggregator	1%
3	No generator export ability and site load to small	Aggregator	1%

3.2.2 Review of Assessment Process Task Importance

The overall results of the importance ratings given by interviewees to each workflow task during the semi-structured interview are shown in Table 8. The importance rating for each task is based on the combined responses from the interviewees, as represented by calculating the first quartile (left line), median (large dot), and third quartile (right line). A notable outcome of these results was the tendency of interviewees to rank all the tasks as being of average to high importance for determining the suitability of a site This is believed to be the result of two factors. Firstly, this is a mature workflow process, refined over time to remove unnecessary tasks that could have received a low importance rating. Secondly, this workflow is directly used by the interviewees as a part of their job, which may cause a bias towards deeming all tasks they undertake as important. This reflects the previously noted social desirability bias limitation of Likert scales, which can result in skewed results as seen in these outcomes (Furnham, 1986). As the results still show variance in importance between each task, the critical tasks for decision-making can be determined based on the highest rankings. Also, due to all median values being between 3 and 5, this bias can be compensated by shifting the rating system whereby a 3 rating is now deemed the lowest importance (instead of 1), and 4 is medium importance (instead of 3), with 5 remaining as highest importance.

Table 8 - Ranked List of all Workflow Tasks with Importance Ratings (1 = Low, 5 = High)

Stage & Task	Task Description	Importance Rating 1 2 3 4 5
1.4	For generator DSR - do they know the size of their generator(s)?	
2.2	For turndown DSR - obtain and analysis detailed list of appliances	
2.3	For generator DSR - obtain detailed information about size and status	
3.3	Organising ENA applications (if needed)	
3.6	Configure control system	
3.7	Perform G59 witness tests (if needed)	
3.8	Site training and user acceptance testing	
3.10	Go live and handover to operations	
1.7	For generator DSR - are they connected in parallel with the grid?	
1.8	For generator DSR - do they have a G59 connection?	
1.10	For turndown DSR - do they have HVAC including chillers, AHU's etc.?	
1.11	For turndown DSR - do they have large fans and/or pumps?	
2.1	Obtain and analysis half hourly electricity usage data	
2.6	Delivery team confirms site potential by either a site or phone survey	
2.7	Undertake contract negotiations	
2.9	Gain client sign-off of contract	
3.5	Arranging subcontractors	
3.9	Spot testing to confirm real DSR potential	
1.6	For generator DSR - are they tested regularly?	
1.9	For turndown DSR - do they have a Building Management System?	
1.12	Discussing programme options	
2.4	For generator DSR - confirm if any existing ability to export	
2.5	Gain initial agreement with client to continue assessment	
3.1	Project Manager and Project Engineer assigned	
3.2	Handover from sales to engineering team	
1.2	Getting in contact with the right person in the company	
1.13	Discussing potential impact to existing operations	
3.4	Arranging equipment purchases	
1.3	Explaining Demand Response and checking to see if interested	
2.8	Create project plan and framework	
1.5	For generator DSR - do they know how old are they?	
1.14	Offering free surveys, monitoring equipment and setup	
1.1	Prequalify potential site/client before making contact	

Looking at the eight highest-ranking tasks in Table 8, five relate to stage 3 and are likely to have received this ranking due to being critical 'must do' tasks for enabling DSR once the contract in stage 2 has been signed. The other three highest rated tasks focus on obtaining asset size information to determine the DSR potential of the site, and reflect the importance placed on this information by the aggregator for deciding client suitability. At the bottom end of the list, the three lowest ranking tasks are all from stage 1, but are all different in purpose and therefore there is no common pattern behind their low importance beyond all belonging to the first stage. The detailed review of each stage in the following sections 3.2.3 to 3.2.5 aims to gain further understanding of the ratings provided based on insights from the semi-structured interviews.

3.2.3 Detailed Review of Process Stage 1 - Initial Client Contact

The results of the semi-structured interview for stage 1 (Initial Client Contact) are shown in Table 9. The table lists the workflow tasks and the importance rating given by the interviewees. The stage 1 results show that the most important criterion for determining client suitability relies on finding out what assets are available for DSR. This is not unexpected as the ability to participate in DSR is reliant on having assets that can be used. The highest rated task is 1.4 in the workflow, which aims to determine the size of any onsite generators. The importance given to this task reflects the UK's focus on this form of DSR, with over 76% of DSR being provided via onsite generators (Grünewald & Torriti, 2013; The Energyst, 2016). The next highest rating is for tasks 1.7 and 1.8, which focuses on trying to determine if the site has the ability to export any spare generator capacity. As outlined in section 2.1.2, the ability to export in the UK effectively means that the full potential of the generator can be utilised for DSR as any excess generation not used by the site during the DSR event is fed into the local network. The importance of exportability was reflected in many of the interviewee's responses, including SJ2 response "*It's great if they have an existing export agreement as we can often double the site's potential as the generators are normally way oversized for the site's normal electricity usage*". As onsite generators are primarily for backup purposes, they are normally sized for peak site demand. This often means that during a DSR event, the site demand can be significantly less than that generator's capacity. If the site can export the spare capacity, then this results in increasing the site's DSR potential and thereby reflects the high importance placed on these tasks. The other equally high importance tasks of 1.10 and 1.11 focus on understanding the site's DSR turndown potential. The high rating of these tasks shows that Turndown DSR is important, and reflects this stage's overall focus of trying to determine the site's DSR potential when assessing whether to progress the client's site to stage 2 (Site Assessment).

Table 9 - Stage 1 - Initial Client Contact Workflow Tasks with Importance Ratings (1 = Low, 5 = High)

Stage 1 Workflow Tasks		Importance Rating 1 2 3 4 5
1.1	Prequalify potential site/client before making contact	
1.2	Getting in contact with the right person in the company	
1.3	Explaining Demand Response and checking to see if interested	
1.4	For generator DSR - do they know the size of their generator(s)?	
1.5	For generator DSR - do they know how old are they?	
1.6	For generator DSR - are they tested regularly?	
1.7	For generator DSR - are they connected in parallel with the grid?	
1.8	For generator DSR - do they have a G59 connection?	
1.9	For turndown DSR - do they have a Building Management System?	
1.10	For turndown DSR - do they have HVAC including chillers, AHU's etc.?	
1.11	For turndown DSR - do they have large fans and/or pumps?	
1.12	Discussing programme options	
1.13	Discussing potential impact to existing operations	
1.14	Offering free surveys, monitoring equipment and setup	

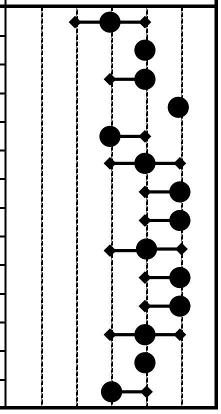

In contrast, comparatively less importance is attributed to tasks that relate to more general assessment areas. The first lower ranking, task 1.1, aims to prequalify a client to save time by only contacting potential clients that are likely to have DSR potential. Interviewee SS2 notes as an example that "*I know from experience that cold store warehouses often state how many pallets they can hold on their websites, so I check and if they have only 10,000 pallet storage then I don't bother as the potential is too low, if they have 100,000 then I contact them*". The low rating for this task is probably due to the difficulty in knowing what potential a site has solely based only on publicly accessible information. As noted by interviewee SJ1: *"I try to find out about potential clients before making contact, but this can be hard, and I tend to just contact them and see if I can get any interest*". The next low rating task 1.5 aims to understand the age of any site generators. Based on interviewees' general comments, task 1.5 is given this low rating as the site contact person often does not know the age of the site's generators, as well as generator age not always providing a good indication of suitability as an old, yet well maintained generator can be more usable than a new, poorly maintained, generator. The importance of maintenance is reflected in the higher rating given to task 1.6, which checks to see if the site's generators are tested regularly, as a better indication of condition.

Table 10 provides a ranked list of the reasons why sites were lost during stage 1 representing, as per Figure 8, the 19% of sites that do not progress to stage 2 (Site Assessment). The list was created based on analysis of KiWi Power's Salesforce database, and is ranked using the Recorded Frequency column, which provides a percentage value for how often the reason occurs during this stage. The Perceived Likelihood rating for each reason is based on the combined responses from the interviews, and shows the first quartile (left line), median (large dot), and third quartile (right line). The results show that there is a contrast between what the interviewees perceived and what the actual reasons were for clients being lost.

Table 10 - Stage 1 - Reasons for Site Loss during Initial Client Contact including recorded frequency and the median and interquartile range of all interviewee perceived likelihood of occurring (1 = Low, 5 = High)

Site Lost Reasons	Recorded Frequency	Perceived Likelihood 1 2 3 4 5
1 - Not Interested in DSR	37%	
2 - No assets or generators large enough for DSR	19%	
3 - Already with another DSR provider	18%	
4 - Assets or generators deemed unsuitable for DSR	14%	
5 - Not Interested due to impact concerns	7%	
6 - Not Interested due to insufficient financial returns	5%	

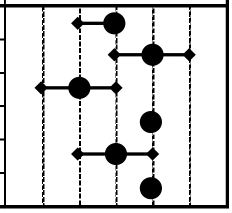

The biggest contrast is seen in the last reason, 6, which is based on clients not being interested due to insufficient financial returns. This had the lowest frequency of occurrence on the sales workflow system, conversely the interviewees ranked this as one of the higher reasons for causing lost sites. Follow-up interview questioning reinforced this result based on the initial financial returns element of discussions with a site often resulted in loss of interest, as the comparatively low returns in relation to overall operating budgets became apparent, with interviewee SS1 noting '*The main reason I have for clients not going ahead is no budget or not financially worth it.'*. The interviewees' perception that clients lose interest due to low returns could help explain why the highest-ranking reason was 'Not interested in DSR'. This high ranking could be the result of it being used as a generic reason where, for example, there is uncertainty about the site's actual specific loss reason, or there is more than one reason, or the KiWi Power staff have not spent the time recording an appropriate reason once the site was deemed lost.

Reasons 2 and 4 for site loss represent the outcomes of this stage's focus on determining if the sites have assets suitable for DSR potential. The interviewees perceived these reasons as being highly likely to cause a site to be lost, which matches the actual frequency of occurrence. Reason 3's middle-ranking demonstrates that DSR in the UK is becoming more common, with a growing base of participating sites. The low-ranking of reason 5, which focuses on concerns about the impact of DSR on a site contradicts the existing research on barriers to uptake as reviewed in section 2.1.4.1 that points to impact concerns being a major reason for rejecting DSR (Olsthoorn et al., 2015; The Energyst, 2016). The reason for this result could be that the interviewees (as they reported) try to avoid discussing impacts during this sales stage, or due to sites having already been lost due to reasons 2 and 4, which occur sooner in the task flow.

3.2.4 Detailed Review of Process Stage 2 - Site Assessment

The results for stage 2 (Table 11) show higher general task importance ratings with less variance than seen in stage 1. The two highest rated tasks 2.2 and 2.3 show, like stage 1, a focus on obtaining detailed information about the site's DSR potential. The importance attributed to these tasks is likely due to the need for an accurate DSR potential assessment for the contract between KiWi Power and the client. The interviewees commented on these tasks as critical for assessing the site's potential, and that it can be difficult and time-consuming to obtain all the detailed records about a site's half-hourly electricity usage data, appliance specifications, wiring diagrams and usage patterns. While task 2.6 is deemed slightly less important, it was remarked upon by interviewees as often being the first time that one of KiWi Power's technicians gets involved with the site. The technicians' task is to confirm the assessment work outcomes with the site's technical knowledge owners by either a phone call or via a site visit. Site visits are avoided unless necessary, as explained by technicians TS1: *"A call is normally done first to see if that can provide enough technical information to confirm the sales team's assessment, if it can't due to complex setup or lack of technical people onsite then we do a site visit, they take more time and money, but about 50% of the time a site visit is needed to make sure"*.

Table 11 - Stage 2 - Site Assessment Workflow Tasks with Importance Ratings (1 = Low, 5 = High)

Stage 2 Workflow Tasks		Importance Rating 1 2 3 4 5
2.1	Obtain and analysis half hourly electricity usage data	
2.2	For turndown DSR - obtain and analysis detailed list of appliances	
2.3	For generator DSR - obtain detailed information about size and status	
2.4	For generator DSR - confirm if any existing ability to export	
2.5	Gain initial agreement with client to continue assessment	
2.6	Delivery team confirms site potential by either a site or phone survey	
2.7	Undertake contract negotiations	
2.8	Create project plan and framework	
2.9	Gain client sign-off of contract	

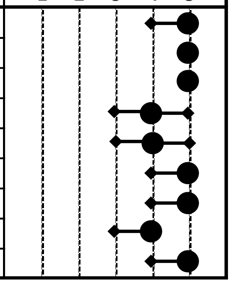

The medium importance ratings given to tasks 2.4, 2.5, and 2.8 mean that they are of comparatively lower importance at this stage. All three relate to non-critical tasks that potentially could be missed without impacting the completion of stage 2. However, even with the lower ratings, these tasks are clearly still seen as important. Task 2.4 aims to confirm existing generator export agreements which, as previously discussed, can lead to doubling the site's DSR potential. However, it is not critical to confirm any existing ability to export before progressing as explained by an interviewee SJ1: "*existing export is nice to have, but normally they don't have it and if site load is good enough we will continue anyway while also trying to obtain an export agreement from their Distributed Network Operator*". The ability to get an export agreement can be difficult though if there are network restrictions, as previously outlined in section 2.1.2. General interviewee comments suggest task 2.5 to be optional, depending on the relationship with the site's client and task 2.8, which deals with project planning, was deemed an administrative task that did not normally impact this stage's overall progress.

Table 12 provides the list of reasons for which 35% of sites (28% of overall sites) were lost during stage 2, as per Figure 8. Like stage 1, the perceived likelihood of reasons for site losses occurring is quite different to the actual recorded frequency of occurrence. The highest-ranking reason at 40% is like the high-ranking reason identified for stage 1, and covers general loss of interest in DSR. As discussed in assessing the reasons for loss at stage 1, its occurrence might be artificially high due to it potentially being a default reason that is used when the real reason is unclear. The next three highest reasons 2, 3, and 4 all have low ratings for perceived likelihood of occurring. There is limited material from the interviewees which explains why this is occurring beyond speculation that reason 2 might relate to company pride given one interviewee SS1 commenting: "*we're normally winning sites from other DSR providers, not losing to them!*".

Table 12 - Stage 2 - Reasons for Site Loss During Site Assessment including recorded frequency and interviewee perceived likelihood of occurring rating (1 = Low, 5 = High)

Site Lost Reasons	Recorded Frequency	Perceived Likelihood 1 2 3 4 5
1 - Lost interest in pursuing DSR	40%	
2 - Choose a different DSR provider	13%	
3 - Lost interest due to internal priorities	10%	
4 - Lost interest due to impact concerns	8%	
5 - Upgrade costs to high for enabling DSR	8%	
6 - Assets deemed technically unsuitable for DSR	7%	
7 - Lost interest due to insufficient financial returns	6%	
8 - DSR potential to low for aggregator to pursue	4%	
9 - Missed deadlines for specific DSR programme	4%	

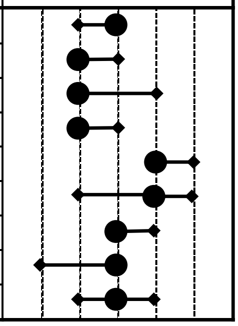

While reasons 4 and 5 have the same recorded occurrence frequencies, each has very contrasting perceived likelihood ratings. Reason 5's high perceived rating matches interviewee comments about enablement costs often causing sites to be lost, as noted by interviewee SJ3: '*Cost of upgrades for G59 connections is the main reason I think we lose sites*'. This is due to backup generators often causing site electricity fluctuations when used, which is acceptable for one-off power cut emergencies, but not for the regular generator usage required for DSR and therefore requires additional equipment to reduce operational impacts.

The two low occurring reasons 7 and 8 relate to lost clients resulting from insufficient financial returns. This is similar to what was seen in the stage 1 reasons for client site losses. Like stage 1, this low occurrence could be related to the workflow task order of stage 2 causing other reasons for loss to occur first, with technical analysis then occurring before financial elements are discussed. Additionally, reason 5 (upgrade costs too high) is related to financial returns on the basis that if returns were high enough, then the costs could be justified.

3.2.5 Detailed Review of Process Stage 3 - Site Enablement and Go Live

The final stage of the sales process deals with the client's site enablement and go live tasks, shifting the focus from assessment to technical delivery. The technical nature of this stage is reflected in most tasks receiving a very high importance rating by both sales and technical interviewees, as seen in Table 13. General interviewee feedback on the high ratings indicated that most of the tasks are essential to enabling the site and, by this stage, contracts had been signed so everyone was motivated to get the site live and earn DSR revenues.

Table 13 - Stage 3 - Site Enablement and Go Live Workflow Tasks with Importance Ratings (1 = Low, 5 = High)

Stage 3 Workflow Tasks	
3.1	Project Manager and Project Engineer assigned
3.2	Handover from sales to engineering team
3.3	Organising ENA applications (if needed)
3.4	Arranging equipment purchases
3.5	Arranging subcontractors
3.6	Configure control system
3.7	Perform G59 witness tests (if needed)
3.8	Site training and user acceptance testing
3.9	Spot testing to confirm real DSR potential
3.10	Go live and handover to operations

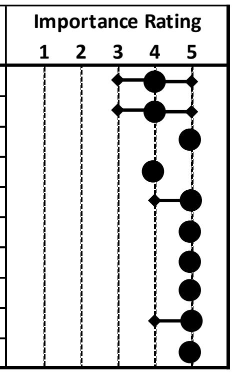

The tasks 3.1 and 3.2 both received comparatively lower ratings, which appears to reflect their softer nature. They relate to organisational tasks without binary outcomes. A project can progress for a while without the presence of an officially assigned project manager, for example, yet configuration or testing is either completed or not. The only contrast to this is task 3.4, which relates to arranging equipment purchases and received a slightly lower rating compared to the other high rating binary type tasks. This could be explained by the response interviewee TJ1 gave when asked about this lower rating: "*equipment is important, however for most of the sites we can use our own control devices which we normally have a good stock of*". This response potentially implies that having the equipment in stock removes a risk during site enablement, and therefore lowers its importance.

This stage has the lowest level of losses at only 18% (6% of overall sites), as per Figure 8. This low level is attributable to sites at this stage having already gone through a detailed assessment, which has filtered out most unsuitable sites. The reasons for the losses experienced during this stage also reflect the change from assessment to technical tasks, with the highest-ranking reason being related to technical issues (as per Table 14). Due to changes in the workflow database, many sites marked as lost at this stage did not have the reason recorded. Therefore, it required asking the interviewees if they knew why each site was lost to gain a full record of reasons for site losses. Not having a full list of reasons for loss during stage 3 at the start of the interviews meant that importance ratings could not be obtained.

Table 14 - Stage 3 - Reasons for Site Loss during Enablement and Go Live including recorded frequency

Site Lost Reasons	Recorded Frequency
1 - Assets deemed technically unsuitable for DSR	31%
2 - Client decided to stop progressing with DSR	16%
3 - Building Management System to old for integration	14%
4 - Turndown reduction to small once tested	14%
5 - Missed deadlines for specific DSR programme	11%
6 - Enablement costs deemed to high	7%
7 - No generator export ability and site load to small	7%

Many reasons for losses (namely 1, 3, 6 and 7) appear to cover areas that the previous stages' assessment tasks were meant to capture. The TS2 interviewee provided insight into the reasons 1 and 3 were rated the highest. In relation to reason 1, technical suitability, the interviewee commented: "*generators are the biggest problem, everything seems fine on paper but once you start working on it you realise it is too old or badly maintained or has some other issue, we now try to focus on this more during the surveys so that we don't find out at this stage*" and thereby implying that more focus is now placed on this during the stage 2 site survey. In relation to reason 3, BMS, the TS1 interviewee proposed: "*this was probably due to one client we had with multiple sites, turns out all the BMS's were ancient so wasn't possible to use them. We find this from time to time, generally if the BMS is older than 10 years it can get tricky to use and is flagged as a risk*".

3.2.6 Minimum DSR kW Potential Requirements

A specific aim of this research was to determine what the minimum kW DSR potential requirements are for KiWi Power to decide if a site is suitable. During the interviews, the participants were all asked what they considered to be the minimum DSR turndown or generator potential required for a site to progress, and why they chose this amount. Figure 9 represents the distribution of interviewees' minimum DSR kW potential requirements for new sites as a box plot, with the whisker representing the minimum and maximum values. For Turndown DSR, the interviewees differentiated between an individual asset's minimum and overall site minimum on the basis that it would not be worth dealing with a site if only one asset at a site met the minimum requirements.

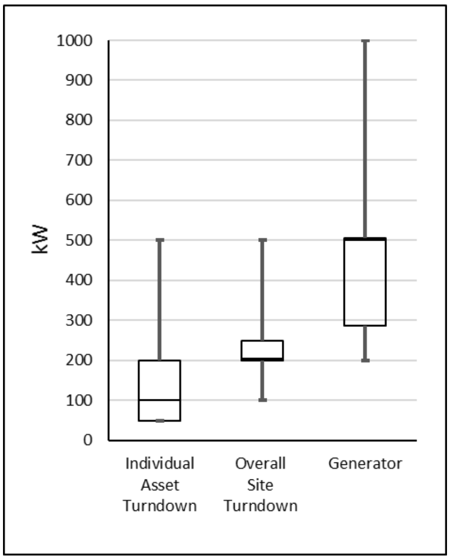

Figure 9 - Distribution of Interviewees Minimum DSR kW Potential Requirements for New Sites

The high variation of minimum DSR potential requirements between and within each category demonstrates a wide range of perspectives from the KiWi Power employees. Looking first at the variance between categories and based on the median values, a typical turndown site would require at least 200kW of turndown potential from two to four assets, each providing between 50kW to 100kW, to be considered by the aggregator as suitable for DSR. In contrast, the median response for the generator category indicates a site would need to provide at least 500kW of generator-based DSR to be deemed suitable. There is no clear reason for this variation across categories beyond the interviewees suggesting it could be attributed to DSR programme revenues, and prior experience. Lower thresholds for turndown suitability were attributed to assets having the potential for being used in the traditionally higher revenue FFR programmes, thereby potentially being profitable despite a smaller DSR potential when compared to generators. Generators cannot normally meet the FFR response time requirements, and therefore require a higher DSR potential to be profitable when using the traditionally lower revenue STOR programme. This might reflect that at the time of the interviews, FFR was paying 2-4 times more revenue than STOR. A limitation of these interviewee responses is that the perception of profitability is based on what the interviewees subjectively consider a profitable kW level, rather than any objective data about actual profitability. Yet interviewee responses about profitability were still sought as the available data from KiWi Power did not include access to any long-term or verifiable analysis by KiWi Power of its required minimum kW levels to ensure a profit. It also appears that prior experience also influences the difference, with interviewees remarking that for turndown it is hard to find large assets, hence accepting that lower levels are more realistic. In relation to generators, interviewee SI1 noted that *"I don't look at any site that has less than 1MW of generator capacity, I know there are tons of sites with large generators out there so it's not worth my time dealing with smaller ones*". This reflects the section 2.1.2 findings that there could potentially be 20GW of backup generator capacity in the UK, of which only approximately 1GW is currently used for DSR, leaving a significant amount of opportunity that supports this interviewee's statement.

Figure 9 also shows high variation within each DSR category. A few of the interviewees provided minimums that they said were based on KiWi Power's 'guidelines'. However, the guideline minimums for each category the interviewees referenced varied, implying that part of the variation is due to confusion about the company's official, mandated minimum amounts. The reason for this was discovered through further questioning about the official guidelines, which revealed that there is no formally recorded amount and what interviewees may deem as guidelines were based on what they were verbally told during induction training, and therefore variable depending on the trainer. A correlation of the interviewees' minimum values against years of experience at KiWi Power provides moderate R-values of -0.50 for individual asset turndown, -0.59 for overall site turndown, and -0.58 for site Generator capacity. Although not a strong relationship, it does indicate that longer experience results in lower minimums. This finding is supported by interview responses. Interviewees with the shortest experience (less than one year) provided some responses about it being hard to know if low DSR potential sites will be worth pursuing, with SJ4 interviewee stating in relation to individual asset turndown: "*I say around 200kW, I think we can go lower but it's hard to know if it will be profitable or not, so I tend to avoid assets with anything less*". These results suggest that the level of uncertainty about the suitability of a low DSR potential site will lower with lesser experience.

3.3 Chapter Conclusion

This chapter has examined the aggregator's role in the adoption of DSR by analysing KiWi Power's workflow system and interviewing its employees. The current complexity and requirements of DSR programmes in the UK means that aggregators are key to helping increase uptake. Therefore, it is important to understand how aggregators operate to help improve uptake of DSR. Studying aggregators also provides valuable information for identifying which DSR commercial barriers require further research. Based on this chapter's findings, only 36% of new end user sites being assessed are likely to go live with DSR. The decision not to continue is made by the client 72% of the time, with the primary reason being loss of interest.

The findings from analysis of KiWi Power's Salesforce workflow system showed that during the first initial assessment stage, the key criteria for determining DSR suitability are based on a high-level identification of the suitable assets for DSR usage. This focused firstly on understanding if the site includes any onsite generators, and secondly on whether the site had any large electrical assets (like HVAC) that could be temporarily turned off. The outcomes of this stage resulted in 19% of sites being lost, with the main reason being 'loss of interest' followed by insufficient DSR potential or unsuitable assets. There was a noticeable discrepancy between interviewee responses, with KiWi Power staff believing that insufficient financial returns are the most likely reason to lose a site, conversely the data from the workflow system shows that insufficient financial returns were actually the lowest recorded reason for lost sites.

The sites that progressed to KiWi Power's second assessment stage underwent a more detailed assessment that involved in-depth analysis of the site's electricity usage and potential DSR assets before agreeing on contractual terms. Findings from this stage showed that most tasks were deemed important for determining site suitability, with the detailed DSR potential analysis tasks having the highest importance rating. This stage experienced the highest level of site losses at 35% with the main reason being sites losing interest in pursuing DSR. A very similar result to stage one, and potentially attributable to this category being used as a generic reason when sites are lost.

The third and final stage of the KiWi Power process involved the site enablement and go live tasks. Almost all tasks in this stage were deemed very important due to most of them being essential for enabling a site for DSR. This stage had the lowest level of site losses at 18%. This reflects that the assessment processes at previous stages have already removed most unsuitable sites. It was noted by interviewees that many of the current reasons for sites being lost have now been addressed by adding additional assessment tasks into the earlier stages. Therefore, it is anticipated that this level will continue to lower for KiWi Power in the future, as the aggregator improves assessment processes to catch any issues before reaching this stage.

The final element of this research was to determine what an aggregator's minimum DSR potential requirements are when deciding site suitability. The results showed that a typical turndown site would require at least 200kW of turndown potential, with a minimum of two assets each providing at least 100kW. Yet, a generator-based DSR site would require at least 500kW. The difference is thought to be the result of DSR programmes at the time of the interviews providing greater returns for turndown sites, and reflect asset sizes, with the potential from turndown assets being smaller than from generators. There is also a moderate correlation that showed that employees with more experience are more likely to accept sites with lower potential.

This chapter's findings highlight three key areas that need addressing to improve DSR uptake. Firstly, understanding the specific DSR potential of a site's assets is the highest priority during the aggregator's assessment process. This means that sites could improve their chances of achieving suitability for DSR by ensuring they understand their electricity usage profiles, and which assets have operational flexibility. Secondly, the main reason for sites not utilising DSR is due to a lack of interest, which is not related to impact or financial concerns. This implies that further end user exposure is required from governments, TSOs and aggregators to drive interest in DSR, and that research is needed into understanding the specific reasons why interest is lost or not generated. Thirdly, the minimum turndown potential of between 50KW to 100kW per asset and 200kW for a site highlights that the financial returns for DSR are insufficient to drive widespread uptake. This will need to be addressed by reducing the costs of enablement and/or increasing DSR programme returns. On the assumption that DSR programmes will not be increasing returns, the focus needs to be placed on understanding how to reduce enablement costs, both through improving processes and lowering technical solution costs. The implications of these findings are further reflected on in the chapter 6 thesis conclusion.

4 A Comparative Analysis of Demand Side Response Estimation Methods

This chapter addresses the second research objective *'To perform a comparison of the outcome uncertainty in DSR potential estimation methods, evaluated against the level of informational requirements of those methods'.* The findings from chapter 3 on KiWi Power's new client assessment process showed that the key criteria for determining suitability focus on understanding the DSR potential of the site's assets. As aggregators normally work on a shared revenue model, accurately determining the potential of a site is important as incorrect assessments can have financial and reputational impacts. The detail is also required to ensure the aggregator will be capable of meeting any contracted levels of reduction required by DSR programmes (as outlined in section 2.1.1), with the threat of penalties applying if missed. As an example, the STOR programme requires participants to provide a guaranteed kW reduction amount for set periods of time of up to 14 hours per day (National Grid, 2016c). If STOR participants underdeliver by more than 5%, then financial penalties are applied and progressively increased, with the potential for ultimately removing non-performing participants from the programme if they fail in meeting guaranteed turn down levels too many times. The severity of penalties will vary by country and DSR programme. For example, the American San Diego Gas & Electric programme has a low severity based on payments being reduced proportionally to the contracted amount delivered (SDG&E, 2016). Conversely the Spanish programme is very strict with exclusion if the site fails to meet their obligations twice (SEDC, 2017). This means that correctly determining the DSR potential of a site is important for appraising its suitability for DSR.

While an aggregator can perform site surveys to gain a detailed understanding of a site's DSR potential, surveys have a time and cost impact. As a result, site surveys are only undertaken by KiWi Power after first deciding that the site is potentially suitable for DSR based on an initial desktop assessment. However, performing a desktop assessment to determine a site's potential is often difficult as detailed usage information (from sub-meters for example) about the individual electrical assets that are being assessed for DSR is normally unavailable (Merry, 2017). Instead, the only information normally available is the site's overall electricity usage over a year, as recorded in half-hourly (UK standard practice) or similar intervals by the site's utility supplier. Half-hourly information will provide a usage profile that can be used for estimating the site's DSR potential and suitability if all electricity demand from the grid is reduced by either using backup generators or turning off all assets. For buildings that can only turndown a limited subset of assets for DSR, a building level profile is unable to provide the individual assets' usage patterns needed to understand their suitability for DSR. To gain this necessary level of detail requires additional analysis, to try and determine what proportion of the site's usage is represented by individual assets.

The methods available for DSR estimation of individual assets are very limited, with existing research on this topic finding only 3 methods, as previously outlined in section 2.3.2 and described in detail in section 4.1.2. While the number of public methods is limited, there are other proprietary commercially developed estimation analysis methods that have not been published. KiWi Power, for example, has supplied one such method in association with this research. This aggregator adopts two complementary approaches when performing site asset assessment for DSR. The first approach assumes that the asset will work at a set level all year. To help reduce the uncertainty of this estimation a second approach is also used which analyses the building's overall electricity records for a year to create a baseload usage amount for 95% of the time. The aggregator then takes a proportion of this 95% to represent the asset usage. Using the baseload value reduces uncertainty, as at least this amount of electricity is being used 95% of the time. Therefore, taking a proportion of it prevents overestimating the asset's potential usage. The major limitation of both approaches is the assumed consistent usage of the asset across the year, which the aggregator recognises, yet still uses the method for the purposes of enabling at least an initial understanding of anticipated DSR potential before deciding on any further investigations.

The issue that faces aggregators and anyone trying to perform DSR estimations using these methods is knowing which one to use, and how the methods compare in terms of uncertainty and cost to undertake. Therefore, the aim of this chapter is to address the second research objective by assessing uncertainty levels in current non-domestic DSR potential estimation methods based on the input requirements. By demonstrating the uncertainty levels and costs of DSR estimation methods, this research hopes to increase the potential for usage of DSR by businesses that are currently excluded due to risk aversion arising from an absence of information about estimation uncertainty levels. This research is undertaken by examining and then applying four DSR estimation methods to two UK hotels, as described in section 4.1. Section 4.2 then sets out the research results and discusses these findings. Section 4.3 concludes by summarising the implications of this research.

4.1 Review of Demand Side Response Estimation Methods

Four DSR potential estimation methods were applied to two UK hotels (each containing approximately 200 rooms) to evaluate outcome uncertainty against the level of information required for estimation. The four methods are: asset assessment; baseline comparison; historical event analysis; and building energy modelling. These methods are used as part of an initial desktop assessment to determine the potential DSR of a building or business. The assessment provides a decision on whether further assessment, including a site survey, is undertaken to decide if the inclusion of the site in a DSR programme is valid. All methods estimate the half-hourly kW usage profile of electrical assets over a one-year period to assess if sufficient DSR potential exists.

The comparison will be undertaken by using each estimation method to generate a DSR asset usage estimation dataset that consists of 365 rows by 48 columns, which represents the half-hourly average kW usage values of the asset over one year. The estimations will be performed on the two hotels' main HVAC chiller assets that as outlined in section 2.1.3 are deemed suitable for DSR due to limited short-term impact on users when turned off or down. The chillers are large centralised assets that cool water for distribution around each hotel's HVAC system to provide space cooling. The hotel chillers have a maximum rating of 333kW for Hotel 1 and 290kW for Hotel 2. The two hotels have been chosen due to having access to detailed information on each site's overall electricity usage along with sub-metered electricity usage of the chillers during the years 2013 and 2016 for Hotel 1 and 2015 and 2016 for Hotel 2. The selection of years was based on the quality of the data, which resulted in other available years being rejected due to missing multiple months of data recordings. The primary reasons for this missing data arise from communication issues (for example, one site required an aerial upgrade for the sub-meter), issues with the chillers preventing them from running, or a temporary suspension of one hotel from providing DSR due to contractual changes. Once each method has generated DSR asset usage estimation datasets for the chiller assets, these datasets will then be compared against the actual chiller usage as provided by the sub-metered information. This will enable a direct comparison of the DSR estimation method outcomes against actual usage, to understand the level of difference and uncertainty introduced with each method and how this compares to their informational requirements.

To explain how the methods were used and compared this section is divided into five sub-sections. Section 4.1.1 describes the main information sources used. Section 4.1.2 reviews each estimation method in detail, including what information sources are required and the calculation steps performed. Section 4.1.3 outlines the comparison methods used to compare the outcomes of the estimation methods. Section 4.1.4 describes the sensitivity analysis approach used to highlight the influence of input parameter uncertainties on method outcomes. Finally, section 4.1.5 describes the approach used to calculate the cost of using each method.

4.1.1 Information Sources

The information required by each estimation method varies, with a combination of shared and unique input requirements. To avoid repetition of information source descriptions, these are all described in detail below, with Table 15 summarising information usage by method:

- **Site Metered Electricity Usage Records**
    - The site's overall electricity usage records as normally provided by the site's utility supplier for one year.
    - Actual usage recorded at a minimum of hourly intervals, preferably sub-hourly.
    - For the hotel test sites, the information is recorded at half-hourly intervals (as per UK standard practice), and in a dataset of 365 rows (or 366 for leap years) representing days of the year, by 48 columns representing each half-hour period of the day.

- **DSR Asset Information**
    - Information about the electrical appliance assets available for DSR usage at the site. A detailed description of potential assets is provided in section 2.1.3.
    - The primary information required is the maximum kW usage rating of the asset. Other information can include, if relevant, whether the asset is influenced by weather (e.g. cooling systems in summer, heating systems in winter) or has set usage patterns (e.g. turned on between 09:00 to 18:00 each day).
    - The DSR assets used in the hotels are the primary centralised chiller units that are used to cool water for distribution around the hotel HVAC system to provide space cooling. The hotel chillers have a maximum rating of: Hotel 1 – 333kW; and Hotel 2 - 290kW, and are all influenced by hot outside temperatures.

- **Historical DSR Event Outcomes**
    - If the site has previously participated in DSR, then it may be possible to obtain the information on the turndown amounts achieved during each event.

- **Weather Information**
    - The primary weather information used is outside air temperature, which has been recorded on the hour from the nearest weather station to the site. Any single missing values are corrected through interpolation. If more than one value is missing in a sequence, then the values are marked as missing with each estimation method then interpreting how to manage the missing data.
    - The hotel test sites utilise weather information obtained from the UK Met Office's - Integrated Data Archive System (MIDAS) (UK Met Office, 2017a).

- **Site building plans and operational information**
    - The sites' building plans are required to construct the Energy Building Model. The plans need to provide enough detail to be able to ascertain for each hotel room sizes and the u-values of the walls, windows, floors, and roof.
    - Site information is also required on the type and operation of the HVAC system and any other information that is available (including hours of operation, occupancy etc.).

- **Asset Sub-Metered Usage Information**
    - Assets that have sub-meters installed can normally record measurements every minute or faster in the form of kW or kWh. This information provides detail about how the specific asset is utilised.
    - The test hotel sites both have sub-meters installed on the chiller assets that are used to test the estimation methods. The data was recorded every minute by taking a reading of the current kW value at that time.
    - The sub-metered data is not used directly by the estimation methods and is instead used by the comparison methods to determine the difference between the estimated and actual values.

Table 15 - Matrix of Information usage by Method

Information Source	Method 1	Method 2	Method 3	Method 4
DSR Asset Information	Y	Y	Y	Y
Site Metered Electricity Usage Records	Y	Y		
Historical DSR Event Outcomes			Y	
Weather information			Y	Y
Site building plans and operational information				Y

4.1.2 Estimation Methods

4.1.2.1 DSR Estimation Method 1 - Asset Assessment

The asset assessment method is based on a review of current estimation approaches undertaken at KiWi Power. The simplest of the four methods, it is based on using very limited information and has two variations. Both variations have been the primary long-term approaches used at KiWi Power for DSR estimation, and have evolved to their current state based on commercial limitations of analysis time and information availability (with half-hourly data being the traditional information source for assessment).

Information Used

This method uses the following sources of information:

- DSR Asset Information.
- Site Metered Electricity Usage Records.

Calculation Methods

The following steps outline the calculations performed for Method 1 - Asset Assessment:

1. **Variation 1 – Minimum Information**
   1.1. An anticipated set percentage usage amount of the asset is selected based on either a default 50%, or another amount if the assessor has prior knowledge of the type of asset and site.

   1.2. The expected kW usage level of the asset is calculated for each half-hour of a year by multiplying the anticipated percentage usage amount by the maximum rating of the asset, with the resulting values being saved into a DSR asset usage estimation dataset.

2. **Variation 2 – Utilise Baseload Calculation**

   2.1. Using the site's Metered Electricity Usage Records, a baseload value is calculated by obtaining the 5th percentile kW value for each half-hour period of the day based on one year's worth of data as per formula (1) (e.g. for each half-hour period of a day, the 365 daily values for the year are obtained and then ranked before determining the 5th Percentile value).

$$n_{HH} = \left\lceil \frac{P}{100} \times N_{HH} \right\rceil \tag{1}$$

Where:
$n =$ kW value of percentile for selected half-hour
$P =$ Percentile
$N =$ Ordered list of kW values for selected half-hour (sorted from least to greatest)
$HH =$ Selected half-hour

   2.2. A percentage value is then selected that represents how much of the baseload is expected to be used by the asset. This can either be a default 10%, or another amount if the assessor has prior knowledge of the asset type and site.

   2.3. The expected kW usage level of the asset is calculated for every half-hour period in a year by multiplying the anticipated percentage usage amount against the baseload kW value, with the resulting values being saved into a DSR asset usage estimation dataset.

   2.4. If the usage outcome is higher than the maximum usage rating of the assets, then the previous step is re-run with a lower percentage.

4.1.2.2 DSR Estimation Method 2 - Baseline Comparison

The second estimation method utilises clustering techniques to identify DSR opportunities through comparison of each building's different usage profiles over a year. This method works on the basis that a building has different usage profiles throughout the year, and once profile clusters are identified, representative profiles of each cluster can be used to ascertain DSR turndown opportunities based on variance between the profiles. Panapakidis et al. (2014) reviewed a selection of clustering methods for electricity load curve analysis of buildings, and identified that the k-means method offers a balanced approach for finding appropriate clusters suitable for understanding building electricity efficiency opportunities, including for DSR. However, these authors did not actually provide specific DSR estimation outcomes for the test building. Research by Van Wijk et al. (1999) also looked into how to use clustering techniques to identify patterns and trends on multiple timescales (days, weeks, seasons). These authors found that using k-means and then associating the resulting clusters to the different timescales allowed for identification and exploration of usage profiles. Their technique succeeds in identifying weekend vs weekday profiles and other significant periods, such as holidays. These clustering techniques show that building electricity usage normally follows a small set of similar profiles. By identifying these profiles, it is then possible to understand different usage levels, which can potentially be used to derive DSR estimations based on the business category.

Using the profiles to estimate DSR requires informed assumptions about what the profiles represent based on available information about the business. For hotels, the focus of this study, assumptions are informed by information on energy sources related to heating and cooling (gas for heating, electricity for cooling), and industry studies/reports on proportional breakdown of electricity use, which identify that HVAC demand typically accounts for 34% of electricity demand in UK hotels (CIBSE, 2012a). The consistent daily profiles of demand across all days of a week, consistent annual occupancy profiles found in hotels, and the high proportion of HVAC related demand suggest that variation in cluster profiles results from differing HVAC loads. It follows that the profile with the highest demand represents a high level of chiller usage, whilst the profile of lowest demand represents a baseline level of chiller usage.

For a different case, such as an office, where weekday and weekend profiles are likely to be represented in different clusters, a larger optimum set of clusters is likely. Identifying baseline level chiller usage would potentially be more difficult in circumstances where greater variability in demand-related activity is found. Determining what the profiles represent highlights the primary drawback of this method, as it requires assumptions to be made on limited data. Incorrectly assuming what the profiles represent will result in incorrect DSR estimations and therefore, this method needs to be used with caution.

Information Sources

This method uses the following sources of information:

- DSR Asset Information.
- Site Metered Electricity Usage Records.
- Weather Information/Degree-days.

Calculation Method

The following steps outline the calculations performed for Method 2 - Baseline Comparison:

1. The k-means cluster method is used for the baseline comparison (Sayad, 2017). This clustering method works by first selecting how many groups the usage dataset will be clustered into. For each group, a random point within the dataset is selected and deemed the centroid value. Each value in the dataset is assigned to the closest centroid. The mean of the values for each centroid is then calculated. The centroids are then moved to the mean position and the values are reassigned to the now closest centroids. This process is repeated until a pre-defined number of interactions is achieved or the level of centroid position change reaches a set tolerance.

2. The number of clusters for the baseline comparison will vary for each site. One approach for determining the optimum number of k-means clusters is known as the 'elbow' method. This method works by repeating the k-means method using a range of clusters to determine each cluster's percentage of variance. The percentage of variance (dependent variable) is plotted against the number of clusters (independent variable) in order to find the 'elbow' of the curve, which signifies the optimum number of clusters as adding more will have limited benefit in reducing variance (Ketchen & Shook, 1996). Figure 10 provides an example of identified `elbow' for clustering of one hotel's daily electricity usage profiles over one year. The main recognised limitations of the elbow method is its reliance on a manual decision-making process to determine where the elbow sits, and that the chart might not have a recognisable elbow if the line is consistent across the clusters (Ketchen & Shook, 1996). The elbow method calculation is performed by:

   2.1. Calculating the percentage of variance explained for a range of clusters (normally 1-15) using the equation (2) (Imran, 2015).

$$\sum_{k=1}^{K} \sum_{i \epsilon S_k} \sum_{j=1}^{p} \left(x_{ij} - \bar{x}_{kj}\right)^2 \tag{2}$$

Where:

S_k = is the set of observations at the k^{th} cluster

$\bar{x}_{kj}$ = is the j^{th} variable of the cluster centre for the k^{th} cluster

2.2. Create a line chart with markers that shows each cluster's percentage of variance as shown in Figure 10 for Hotel 1 in 2016.

2.3. Determine the elbow based on the chart and record the cluster number.

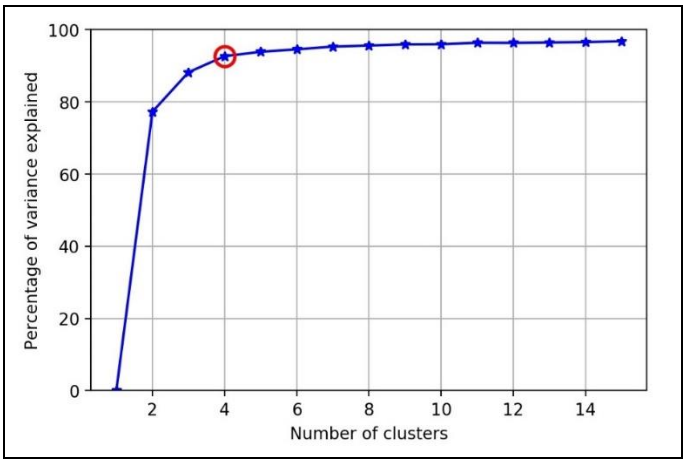

Figure 10 - Example of Cluster Identification using the Elbow Method (with the Elbow being indicated by the red circle)

3. Once the number of clusters to be used has been determined, then the k-means method as shown in equation (3) (Sayad, 2017) can be used to group the Site's Half-Hourly Electricity dataset into similar days. The dataset is then updated with a new column 49 containing a value that represents which cluster each day belongs to.

$$J_n = \sum_{j=1}^{K} \sum_{i=1}^{n} \left(x_i - c_j\right)^2 \qquad (3)$$

Where:
$n =$ Objects being clustered
J_n = Cluster outcome for n value
$K =$ Clusters
c_j = Centroid for cluster *j*
x_i = Object *i*

4. The half-hourly averages in each cluster are then used to generate daily profiles at half-hourly resolution for each cluster of each hotel. Figure 11 provides an example of the daily profiles developed for the four identified clusters of a hotel.

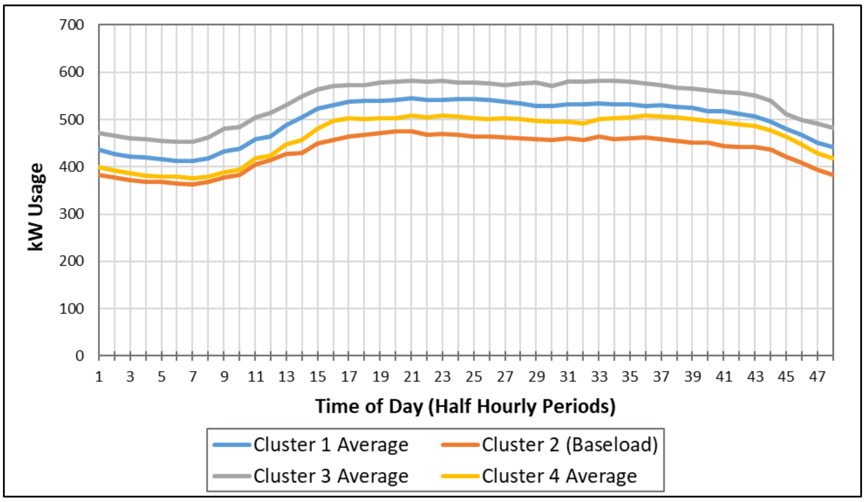

Figure 11 - Example Chart of Clustered Averages

5. The baseline profile is then identified based on the assumption that the profiles represent differences in chiller usage levels. In the context of the UK, chillers are not typically in use during the winter months. Therefore, the baseline is considered as days when the chiller is switched off during the heating season and, as a result, profile cluster 2 in Figure 11 comprises the baseline profile as it has the lowest usage values. The remaining cluster profiles then represent days when the chiller is in use.

6. A new dataset is created that covers all half-hourly periods for one year, and has an additional column identifying which cluster profile is associated with each day of the year. For each day in the dataset, the kW usage levels of the chiller is estimated by the calculating the difference between that day's cluster profile usage value and the baseline value. If a day in the new dataset is associated with the baseline cluster, then the chiller is deemed to be off during this day, so the expected usage is set to 0.

7. The dataset now represents the DSR asset usage estimation dataset of the chiller. The results are then checked to verify that no values are greater than the maximum usage rating of the chiller asset. If there are, then the values are adjusted down to the maximum rating or, if the values are consistently too high, then this method is rejected if the assessor believes the method is providing unrealistic results based on the assessor's (or their colleagues') prior knowledge of customary usage for this type of asset.

4.1.2.3 DSR Estimation Method 3 - Utilise Historical DSR Event Outcomes

If a building has previously participated in DSR, then information gained on the kW amount reduction during each event can be utilised to estimate future performance. Research on this method has traditionally focused on confirming the post DSR event performance of a building by calculating the turndown amount achieved which is deemed as the difference between normal non-DSR building usage and the actual usage during a DSR event (Mathieu et al., 2010). Further research into understanding the expected level of turndown demand using weather-based regression analysis was undertaken by Piette et al. (2011). They showed that the level of turndown potential could be linked to temperature if the DSR assets demonstrate varying levels of usage based on external weather conditions. This DSR estimation method utilises these concepts to identify a predictor that determines the expected turndown amount of historical DSR events. The predictor can then be utilised to determine the expected turndown amount at any time over a one-year period.

Information Sources

This method uses the following sources of information:

- Historical DSR Event Outcomes.
- Weather Information.

Pre-Method Steps – Creation of Sample Event Turndown Results

This method relies on access to historical DSR event outcomes for the building. To provide consistency for testing this method with both hotels, a set of 24 sample DSR events were created using a weighted random sample from 3 years of UK STOR DSR events, as experienced by KiWi Power. The steps undertaken to create the sample events are outlined below:

1. In Excel, a list with the date and times of 77 actual STOR events between 01/09/2012 and 30/04/2017 is compiled based on records obtained from KiWi Power.

2. A count of events per day of the year is created. The percentage probability of occurring per day is then calculated before being converted into the cumulative probability distribution.

3. The cumulative probability distribution represents the likelihood of an event occurring on any given day of the year. This is used to obtain a weighted random sample of 24 days in the year that an event might occur using historical STOR events. The selected days were saved into 'Random Day' column of Table 16 and Table 17.

4. A count of events per half-hour period of the day is created. The percentage probability of occurring per half-hour is then calculated before being converted into the cumulative probability distribution.

5. The cumulative probability distribution is used to find 24 half-hour periods as per step 3, and saved into the 'Random Time' columns of Table 16 and Table 17.

6. The random day and half-hour period for each simulated event is then used to determine the associated actual data and time for each hotel's two dataset years, and then saved into the 'Event Date/Time' columns of Table 16 and Table 17.

7. The event date/time value is then used to obtain the actual sub-metered kW usage values of the DSR asset being assessed, and saved into the 'Event kW Value' columns of Table 16 and Table 17. The simulated events work on the basis that the asset was turned off for one hour (as is the standard KiWi Power practice for these sites) with the recorded average usage for that hour being deemed the turndown amount achieved.

Table 16 - Hotel 1's Simulated DSR Events for Estimation Method 3

Event	Random Day	Random Hour	Simulated Event Date/Time 2013	2013 Event kW Value	Simulated Event Date/Time 2016	2016 Event kW Value
1	18	35	18/01/2013 17:30	46	18/01/2016 17:30	49
2	20	38	20/01/2013 19:00	46	20/01/2016 19:00	52
3	31	22	31/01/2013 11:00	41	31/01/2016 11:00	21
4	42	38	11/02/2013 19:00	19	11/02/2016 19:00	20
5	64	30	05/03/2013 15:00	41	04/03/2016 15:00	11
6	65	35	06/03/2013 17:30	41	05/03/2016 17:30	41
7	78	33	19/03/2013 16:30	105	18/03/2016 16:30	44
8	81	36	22/03/2013 18:00	41	21/03/2016 18:00	15
9	165	21	14/06/2013 10:30	63	13/06/2016 10:30	64
10	205	38	24/07/2013 19:00	67	23/07/2016 19:00	49
11	234	16	22/08/2013 08:00	121	21/08/2016 08:00	91
12	245	35	02/09/2013 17:30	112	01/09/2016 17:30	80
13	248	34	05/09/2013 17:00	212	04/09/2016 17:00	155
14	256	35	13/09/2013 17:30	149	12/09/2016 17:30	90
15	258	22	15/09/2013 11:00	124	14/09/2016 11:00	126
16	272	21	29/09/2013 10:30	117	28/09/2016 10:30	147
17	281	17	08/10/2013 08:30	138	07/10/2016 08:30	113
18	289	38	16/10/2013 19:00	72	15/10/2016 19:00	101
19	292	19	19/10/2013 09:30	39	18/10/2016 09:30	84
20	315	34	11/11/2013 17:00	58	10/11/2016 17:00	73
21	316	33	12/11/2013 16:30	48	11/11/2016 16:30	72
22	320	14	16/11/2013 07:30	47	15/11/2016 07:30	67
23	321	18	17/11/2013 09:00	47	16/11/2016 09:00	67
24	326	22	22/11/2013 11:00	45	21/11/2016 11:00	75

Table 17 - Hotel 2's Simulated DSR Events for Estimation Method 3

Event	Random Day	Random Hour	2013 Event Date/Time	2015 Event kW Value	2016 Event Date/Time	2016 Event kW Value
1	18	35	18/01/2015 17:30	0	18/01/2016 17:30	0
2	20	38	20/01/2015 19:00	0	20/01/2016 19:00	0
3	31	22	31/01/2015 11:00	0	31/01/2016 11:00	0
4	42	38	11/02/2015 19:00	0	11/02/2016 19:00	0
5	64	30	05/03/2015 15:00	0	04/03/2016 15:00	0
6	65	35	06/03/2015 17:30	0	05/03/2016 17:30	0
7	78	33	19/03/2015 16:30	29	18/03/2016 16:30	0
8	81	36	22/03/2015 18:00	0	21/03/2016 18:00	0
9	165	21	14/06/2015 10:30	22	13/06/2016 10:30	0
10	205	38	24/07/2015 19:00	25	23/07/2016 19:00	58
11	234	16	22/08/2015 08:00	31	21/08/2016 08:00	104
12	245	35	02/09/2015 17:30	23	01/09/2016 17:30	101
13	248	34	05/09/2015 17:00	101	04/09/2016 17:00	153
14	256	35	13/09/2015 17:30	38	12/09/2016 17:30	136
15	258	22	15/09/2015 11:00	33	14/09/2016 11:00	126
16	272	21	29/09/2015 10:30	49	28/09/2016 10:30	120
17	281	17	08/10/2015 08:30	65	07/10/2016 08:30	122
18	289	38	16/10/2015 19:00	0	15/10/2016 19:00	126
19	292	19	19/10/2015 09:30	0	18/10/2016 09:30	131
20	315	34	11/11/2015 17:00	0	10/11/2016 17:00	119
21	316	33	12/11/2015 16:30	0	11/11/2016 16:30	113
22	320	14	16/11/2015 07:00	0	15/11/2016 07:00	104
23	321	18	17/11/2015 09:00	0	16/11/2016 09:00	47
24	326	22	22/11/2015 11:00	0	21/11/2016 11:00	47

Calculation Method

The following steps outline the calculations performed for Method 3 - Utilise Historical DSR Event Outcomes:

1. The first step is to determine what variables are available for predicting the event turndown amount. For this example, the variables of Outside Air Temperature, Site Electricity Usage, Half-Hour Period of Day, and Day of Week are used.

2. For each variable, a two-column dataset is created for each year of data with the first column containing the event turndown results, and the second column containing the predicting variable value.

3. Using equation (4) the R-squared/coefficient of determination for each dataset is calculated.

$$R^2 = 1 - \frac{\sum_i (y_i - f_i)^2}{\sum_i (y_i - \bar{y})^2} \quad (4)$$

Where:
R^2 = R-squared/coefficient of determination
y_i = Current value from event data set
$\bar{y}$ = Mean of event data set values
f_i = Predicted value for y_i

4. The R-squared values of each variable used as shown in Table 18 are compared, and the highest value selected as the predictor variable to be used for estimating DSR asset usage. In this case the Outside Air Temperature has the highest values.

Table 18 - Method 3's R-squared Regression Results

Hotel/Year	Time of Day	Day of Week	Site Electricity Usage	Outside Air Temperature
Hotel 1 - 2013	0.003	0.036	0.273	0.722
Hotel 1 - 2016	0.003	0.040	0.087	0.636
Hotel 2 - 2015	0.007	0.017	0.046	0.434
Hotel 2 - 2016	0.019	0.028	0.066	0.447

5. The Outside Air Temperature values for each half-hourly period of the year in conjunction with the predictor's slope and y-intercept are used to calculate the DSR estimation potential for the hotels.

4.1.2.4 DSR Estimation Method 4 - Building Energy Modelling

Building energy modelling provides insight into DSR potential by modelling the energy usage of building assets under different operational and environmental scenarios. Modelling gives insight into the flexibility of asset usage that can then be used for DSR estimation. However, this is very time-consuming in comparison to the previous estimation methods, and requires a very high level of information and specialised skills to complete. Utilising a database of archetypal building models for a building stock can help reduce the modelling burden for DSR, as demonstrated by Yin & Black (2015). The predefined model archetypes can be modified as necessary, but its success is dependent on the maturity of the database of archetypes and level of modification needed to provide results that can be used for DSR estimation. Another issue with energy building models is the 'performance gaps' between model designs and actual performance of completed buildings, which can result in high levels of output uncertainty (Menezes, Cripps, Bouchlaghem, & Buswell, 2012). For this research, the building energy model DSR estimation method utilises the Yin & Black (2015) methodology by creating a building energy model of the hotels using EnergyPlus. The outcome of the simulation includes the expected level of kW required for HVAC per half-hour that will be used for DSR estimation.

Information Used

This method uses the following sources of information:

- Site Building Plans and Construction Cross Sections.
- DSR Asset Information.
- Weather Information.

Calculation Method

The following steps outline the calculations performed for Method 4 - Building Energy Modelling:

1. The building plans for each hotel were used to provide both accurate building dimensions as well as the fabric structure of the building (outlined in Table 19). The building plans are used to create a representative model of the building using the software package 'DesignBuilder' v5.0.2 (DesignBuilder, 2017b), as shown in Figure 12. The DesignBuilder program then utilises the EnergyPlus simulation program (U.S. Department of Energy, 2017) to estimate the building's energy usage over one year at half-hourly intervals.

Table 19 - Build Energy Model Components

Component	Hotel 1 Description	Hotel 2 Description
External Walls	400mm thick wall (formed of stone masonry, brick, glass wool insulation, and plasterboard) total U-Value of 0.289	300mm thick wall (formed of brick, polystyrene insulation, concrete, and plasterboard) total U-Value of 0.351
External Windows	Double glazed (formed of two 3mm panes with a 6mm air gap) total U-Value of 3.365	Double glazed (formed of two 3mm panes with a 6mm air gap) total U-Value of 3.365
Roof	400mm flat roof (formed of asphalt, glass wool insulation, air gap, plasterboard) total U-Value of 0.322	320mm Flat roof (formed of asphalt, glass wool insulation, air gap, plasterboard) total U-Value of 0.346
HVAC System	Fan Coil Unit (4-Pipe), 333kW air-cooled chiller with a cooling set point of 23°C	Fan Coil Unit (4-Pipe), 290kW air-cooled chiller with a cooling set point of 23°C
Property Details	7 storeys, ~21,000 m^2 isolated building located in Bristol, UK.	6 storeys, ~15,000 m^2 isolated building located in London, UK.

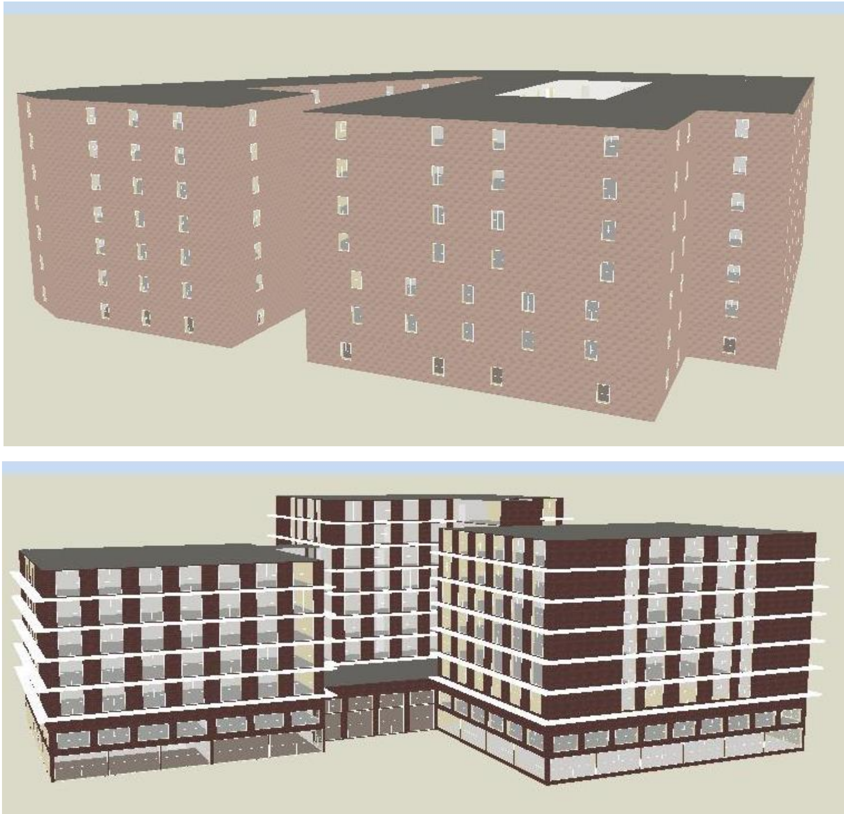

Figure 12 - Building Energy Model Renderings of Hotel 1 and Hotel 2

2. Customised weather files were generated for each hotel for the years 2013 and 2016 and loaded into DesignBuilder. These were created using MIDAS weather data (UK Met Office, 2017a) that was then converted into an EnergyPlus formatted hourly weather data .epw file using the process outlined on the DesignBuilder online help (DesignBuilder, 2017a).

3. Each model's energy usage was then simulated at half-hour intervals for one year using DesignBuilder/EnergyPlus, with the results of the chiller assets electricity usage being extracted to provide the DSR estimation potential for each hotel.

4.1.3 Comparison Method

The selection of appropriate statistical evaluation methods for forecasting is complicated by the many options available. This is highlighted in a paper by De Gooijer & Hyndman (2006), who reviewed 25 years of time series forecasting and started their evaluation section with *'A bewildering array of accuracy measures have been used to evaluate the performance of forecasting methods.'* Therefore, appropriate evaluation measures for this research were determined by first selecting the existing approaches used by previous researchers for DSR estimation and then comparing these against general reviews of forecasting evaluation methods. The review and justification of the selected methods is provided in the following sub-sections. Each of the statistical methods will be applied to the outcomes of the four DSR estimation methods defined in section 4.1.2. The reasons for each method's selection and a detailed overview of its usage are also provided in this section.

4.1.3.1 Evaluation Method 1 - Mean Absolute Percentage Error (MAPE)

The Mean Absolute Percentage Error (MAPE) as defined in equation (5) is the first statistical evaluation method to be utilised. This method has been chosen as De Gooijer & Hyndman (2006) define it as one of the most common measures to use for time series evaluation, and given its general usage across the literature on estimation methods (Aman et al., 2016; Larsen, Pinson, Leimgruber, & Judex, 2015). As well as comprising a standard evaluation method, MAPE also provides the benefits of normalising the outcome in a percentage form for ease of understanding, and the usage of absolute values ensures the resulting error value is not hidden by positive and negative differences cancelling each other out. Its main drawbacks include divisional errors if there are zero values in the dataset, scale sensitivities that result in very high percentages outputs if there are significant differences in the actual and forecasted values, and inability to show forecast bias. While the MAPE output provides a good measure for comparing the overall error levels of different DSR estimation methods its inability to show error is a major limitation when used for DSR forecasting. This is due to DSR programmes normally penalising for underperformance as described in section 2.3.1. This means that an estimation method with a positive bias could have a greater financial impact than one with a negative bias, due to overestimating the site's potential. Therefore, to address this limitation, the Mean Bias Error (MBE) method described in the following 4.1.3.2 will be used in conjunction with the MAPE for assessing the outcomes of the DSR estimation method comparison.

The MAPE method is undertaken by first calculating the difference between the forecasted and actual electricity usage and then dividing this value by the actual usage for each half-hour period of the year. The absolute value of the outcomes is then summed and divided by the number of data points used. The resulting value provides a percentage indication of the level of fit between the forecasted and actual values, with lower being better.

$$MAPE_{year} = \frac{\sum_1^n \left|\frac{A-F}{A}\right|}{n} \tag{5}$$

Where:
F = Forecasted value
A = Actual value
n = Number of values

4.1.3.2 Evaluation Method 2 - Mean Bias Error (MBE)

The Mean Bias Error (MBE) as defined in equation (6) is the second statistical evaluation method chosen. This method has been selected as it is recommended in the ASHRAE guidance 14 'Measurement of energy and demand savings' as providing an easy to interpret percentage measure of the difference between forecasted and actual electricity usage values (ASHRAE, 2014). It is also the method utilised by the Method 4, 'Building Energy Modelling' researchers (Dudley, 2010). The primary benefits of the MBE method comprise an easy to interpret percentage output that shows the overall level of difference between the forecasted and actual values, with lower values indicating a better outcome, and its ability to show if the forecast differences are positively or negatively biased, which helps overcome the MAPE method's inability to show forecast bias. Its main drawback is that the result can be impacted by offsetting errors, which can result in large positive and negative difference values neutralising each other, resulting in a low MBE that does not accurately represent the actual outcomes. The MBE and MAPE methods complement each other through the MAPE's ability to show overall levels of error, and the MBE's ability to indicate if the error levels have a positive or negative bias. As both methods are being applied to year-long forecasts using half-hourly intervals, error outputs are an average of 17,520 differences between actual and estimated values. This will result in reducing the impact of large errors, for example, if the 1,488 half-hour January forecasts were all 200% higher than the actual values with the rest of the year being 10% under, as then the MAPE would be 26% with an MBE of 8%. This could cause DSR programme participation issues if the forecasts alone were used to determine the site's DSR potential due to significant underperformance during January. However, as outlined at the start of this chapter, the DSR estimation methods reviewed are intended for use during the initial desktop suitability assessment of a new site only. If the site passes this stage, then as described in section 3.1.1, additional site surveys and testing should be undertaken before confirming the actual DSR potential that the site would be expected to provide. Therefore, the MAPE and MBE methods provide suitable measures of error for understanding how the DSR estimation methods compare.

The MBE method is undertaken by first calculating the difference between the forecasted and actual electricity usage for each half-hour period of the year. The differences are summed and then divided by the sum of the measure electricity usage over the year. The resulting percentage shows the overall level of difference and bias between the forecasted and actual values.

$$MBE_{year} = \frac{\sum_1^n (F-A)}{\sum_1^n (A)} \tag{6}$$

Where:
F = Forecasted value
A = Actual value
n = Number of values

4.1.3.3 Sub-Meter Data Used for Estimation Evaluation

To use the statistical evaluation methods requires having actual data values to compare the estimated values against. For this research, sub-metered data of the HVAC chillers at the two hotels was extracted from the KiWi Power 'KiWi Operation Management Platform' (KOMP) system. The sub-meter data consists of minute intervals recordings of chillers' current kW usage level at the time of recording. The available sub-metered data was assessed for consistency, which resulted in the datasets for years 2013 and 2016 being selected for Hotel 1 and 2015 and 2016 being selected for Hotel 2.

To utilise the selected datasets required converting them into the same half-hourly format as generated by the DSR estimation methods. To perform this conversion, a Python programme was created that undertook the following steps:

1. The sub-metered chiller kW usage data is obtained in CSV format, and loaded into the Python programme on a per hotel basis and separated by year.

2. Each year of minute interval data is converted into half-hourly values by summing the kW values over each half-hour, and then dividing by 30.

3. During the conversion process, missing values within one half-hour conversion period are addressed by dividing the summed values with the number of available recordings (e.g. if there are only 29 recorded values in one half-hour period, then they are summed and divided by 29). If a conversion period has no record values, then a value is derived by interloping the previous and next conversion outcomes. If there is more than one sequential period of missing values, then these periods are marked as missing data that the MAPE and MBE methods ignore.

4.1.3.4 Evaluation Calculation Method

To calculate the MBE and MAPE for reach DSR estimation outcomes, the following steps were performed:

1. Each DSR estimation method's dataset output was converted from a table of 365 rows (days) by 48 columns (half-hour periods) into a three-column list of: Date, Half-Hour Period, Estimated kW Value.

2. The actual sub-metered usage data was converted from a table of 365 rows (days) by 48 columns (half-hour periods) into a three-column list of: Date, Half-Hour Period, Actual kW Value.

3. Each DSR estimation method's list of estimation values was linked to the associated actual values list.

4. The MBE and MAPE evaluation methods equations as defined in sections 4.1.3.1 and 4.1.3.2 were then applied to each year of data. The resulting MBE and MAPE values for each estimation method and year were recorded.

4.1.4 Uncertainty Analysis Approach

The accuracy of estimation method is an important factor in creating credible, robust DSR portfolios that can meet grid-operator needs. Therefore, appropriate interpretation of uncertainty in inputs used by the proposed methods is critical to DSR estimation. Selecting the right uncertainty analysis approach is important as highlighted by Uusitalo et al. (2015), who evaluated the following approaches and when their usage is appropriate: expert judgement, model emulation, sensitivity analysis, use of multiple models, and statistical approaches. Based on their evaluation, sensitivity analysis was selected for this research as the approach providing the most suitable fit for gauging the level of uncertainty as DSR estimation methods have limited and defined inputs that directly influence an estimation's outputs.

The sensitivity analysis approach is further outlined by Saltelli, Chan, & Scott (2008), who explains it in this way: *"sensitivity analysis studies the relationship between information flowing in and out of the model*". In the case of the estimation methods, this approach is applied by studying the key input variables of each method and how these variables impact the resulting estimated kW output values. To understand the impact of each estimation method's input uncertainty on the DSR estimation, and so give insight as to where more accurate information should be sought, a one-at-a-time local sensitivity analysis test was carried out (Saltelli, Chan, & Scott, 2008). Table 20 provides a summary of the input parameters used, and how these were adjusted from the base values that were defined in section 4.1.2 for each estimation method. As estimation methods 1-3 only have one or two input variables, all inputs for each method are tested during the analysis. In contrast, the detailed modelling approach of method 4 has a wide range of input variables, ranging from building form and structure, to operational schedules of appliances and occupancy profiles. In this instance, it is assumed that the availability of building plans and detailed information on HVAC and lighting infrastructure reduces uncertainty in many of the structural aspects of the model. Menberg, Heo, & Choudhary (2016) identified temperature set points, thermal conductivity, and air infiltration as having significant impact on building energy model results. These three variables are the focus of our analysis for Method 4.

Table 20 - Summary of Estimation Method Sensitivity Analysis Input Parameters

Method	Input Parameter Adjusted	Method Base value	Sensitivity Intervals (including base value in **bold**)
1 (1)	Adjust asset usage percentage by +/- 5 and 10 points	50%	40%, 45%, **50%**, 55%, 60%
1 (2)	Adjust asset percentage usage of baseload value by +/- 2.5 and 5 points	10%	5%, 7.5%, **10%,** 12.5%, 15%
	Adjust baseload percentile by +/- 1 and 2 points	5%	3%, 4%, **5%,** 6%, 7%
2	Adjust number of clusters by +/- 1 cluster	4	3, **4**, 5
3	Adjust number of available existing events by -50%, +50%, +100%	12	6, **12**, 18, 24
4	Adjust cooling set point by +/- 1 and 2 °C	23 °C	21, 22, **23**, 24, 25
	Adjust U-Values of External Walls, Windows, and Roof by +/- 10% and 20% u-values	Hotel 1: Wall 0.289 Window 3.365 Roof 0.346 Hotel 2: Wall 0.351 Window 3.365 Roof 0.322	Hotel 1: Wall - 0.231, 0.260, **0.289**, 0.318, 0.348 Window - 2.692, 3.028, **3.365**, 3.701, 4.038 Roof - 0.277, 0.311, **0.346**, 0.380, 0.415 Hotel 2: Wall - 0.280, 0.316, **0.351**, 0.386, 0.421 Window - 2.692, 3.028, **3.365**, 3.701, 4.038 Roof - 0.258, 0.290, **0.322**, 0.354, 0.386
	Adjust air infiltration levels by +/- 0.1 and 0.2 ac/h	0.7	0.5, 0.6, **0.7**, 0.8, 0.9

4.1.4.1 Sensitivity Analysis Calculation Method

To undertake the sensitivity analysis, the following steps were performed for each estimation method:

1. Each estimation method was run as outlined in section 4.1.2 using the base value inputs, with the resulting MAPE, MBE and MWh output values being recorded into a spreadsheet.

2. Each estimation method was then repeated with one input value being changed as per Table 20, and the output values being recorded into the spreadsheet.

3. Once all input variations for one method had run, the spreadsheet was updated to calculate the following normalised percentage change values:
    a. The input parameter percentage change from the base value.
    b. The MWh output percentage difference from the base value.

4. The information was then summarised into a table and charted for each estimation method.

4.1.5 Determining the Cost of each Estimation Method

The final output of the review of DSR estimation methods is a comparison of each method's estimation errors in relation to its cost to run. This comparison is performed to provide context on usage of each method in a business setting. It enables consideration of the cost/benefit selection of a higher error method that is cheaper or vice-versa. Calculating each method's cost to run in a business setting required estimating the time it would take an experienced user to perform the tasks needed to run the estimation method, and the cost of any external data input requirements. Table 21 provides a summary of the expected time required and external cost (if any) for each informational input. The user time is subjective, by person and business. Therefore, the figures used are estimations based on experience gained through application of these methods for this research and observations of users within KiWi Power. The time value includes both the time taken to obtain information about the building (this covers talking to the building representative to obtain the site's half-hourly electricity usage data and information about the DSR assets) and the time required to format, analyse, and interpret the data. Most external information has no direct cost, as it is obtained for free from the building users or other sources. The only externally sourced information incurring cost is historical weather observations (ECMWF, 2017), which has a fixed yearly fee of £5,000 and has been split into individual usage costs on the assumption of performing 500 assessments per year (£10 per usage).

Table 21 - Summary of Estimation Methods Information Input Costs

Information Input	User Time to obtain/use (minutes)	User Time cost (@ £20 per hour)	External Information Cost	Cost of External Information (@500 uses year)	Total Input and Usage cost
Maximum kW rating of sites DSR assets	30	£10.00	Free	£0	£10.00
Site electricity usage records for 1 year	60	£20.00	Free	£0	£20.00
Previous DSR Event Outcomes	120	£40.00	Free	£0	£40.00
Hourly outdoor weather information for 1 year (ECMWF, 2017)	60	£20.00	£5000 per year	£10	£30.00
Site building plans and operational information	420	£140.00	Free	£0	£140.00

To calculate the total cost of performing each method, the individual costs of gaining data for each input from Table 21 are associated with each method as per Table 22. This table shows the cumulative total running cost of each method, based on the information required. This information combined with the MAPE results from section 4.1.3 enables a comparison of estimation error against method cost to be performed, as shown in section 4.2.3.

Table 22 - Summary of Costs to Perform Each Estimation Method

	Information Usage and Cost per Method				
Information Input & Cost	**1 (1)**	**1 (2)**	**2**	**3**	**4**
Maximum kW rating of sites DSR assets	£10	£10	£10	£10	£10
Site electricity usage records for 1 year		£20	£20		
Previous DSR Event Outcomes				£40	
Hourly outdoor weather information for 1 year				£30	£30
Site building plans and operational information					£140
Total Cost per Method	**£10**	**£30**	**£30**	**£80**	**£180**

4.2 Results of Estimation Method Comparison

The results of applying the four DSR estimation methods to two hotels are reviewed and discussed over three sections. Section 4.2.1 reviews the initial outputs of each method by applying 'base case' values to the input variables, and comparing the estimation error between methods. Section 4.2.2 then reviews the sensitivity analysis results to understand the impact of input variables on the estimation error levels. Section 4.2.3 reviews how the error levels compare against the estimated cost of undertaking each method, to gain an understanding of how cost and error levels correspond.

4.2.1 MAPE and MBE Estimation Method Outcomes

The estimation errors of MAPE and MBE for each estimation method, when using default (base) values for input variables, are given in Table 23. The methods were applied to each hotel over two years to generate a predicted half-hourly kW usage value for their HVAC chillers. The predicted kW values were then compared to the actual kW usage values (as recorded by sub-meters), and MAPE and MBE were calculated for annual estimation errors. The average, minimum, and maximum MAPE and MBE values were then calculated, as shown in Figure 13 (see Table 45 in Appendix D for figure numbers). The MAPE values provide an overall indication of the level of difference between the actual and predicted results. Figure 13 and Table 23 show a range of MAPE estimation errors across the methods, with M1-V1 'Asset Assessment' having the highest average level of error at 159%. In contrast, M3 'Utilise Historical DSR Event Outcomes' had the lowest average level of error at 39%.

The MBE values indicate the direction of error between the actual and prediction values, with positive and negative results indicating overestimation and underestimation respectively. Figure 13 shows that all methods, except M1-V1, under predict usage levels. As seen with the MAPE result, the M1-V1 outcome also has the highest average MBE value at 150%, which indicates that this method overpredicts the expected usage of the HVAC chiller. In contrast, with an average MBE of -10%, M4 predicts most closely to the annual average DSR potential.

Table 23 - Individual Hotel Summary of Estimation Method Error Levels

Method	Hotel 1 - 2013		Hotel 1 – 2016		Hotel 2 - 2015		Hotel 2 - 2016	
	MAPE	MBE	MAPE	MBE	MAPE	MBE	MAPE	MBE
M1-V1	193%	122%	250%	136%	98%	236%	96%	104%
M1-V2	35%	-46%	59%	-50%	71%	1%	75%	-38%
M2	57%	-41%	59%	-29%	40%	16%	70%	-12%
M3	33%	-15%	40%	-7%	36%	-6%	46%	-18%
M4	58%	-1%	63%	5%	39%	2%	45%	-31%

Abbreviation Key:

M1-V1 = Method 1- Variation 1 - Minimum information using set percentage of asset usage

M1-V2 = Method 1- Variation 2 - Utilise baseload calculation with set usage percentage

M2 = Method 2 - Baseline comparison using cluster analysis

M3 = Method 3 - Regression analysis utilising historical DSR event outcomes

M4 = Method 4 - Building energy modelling

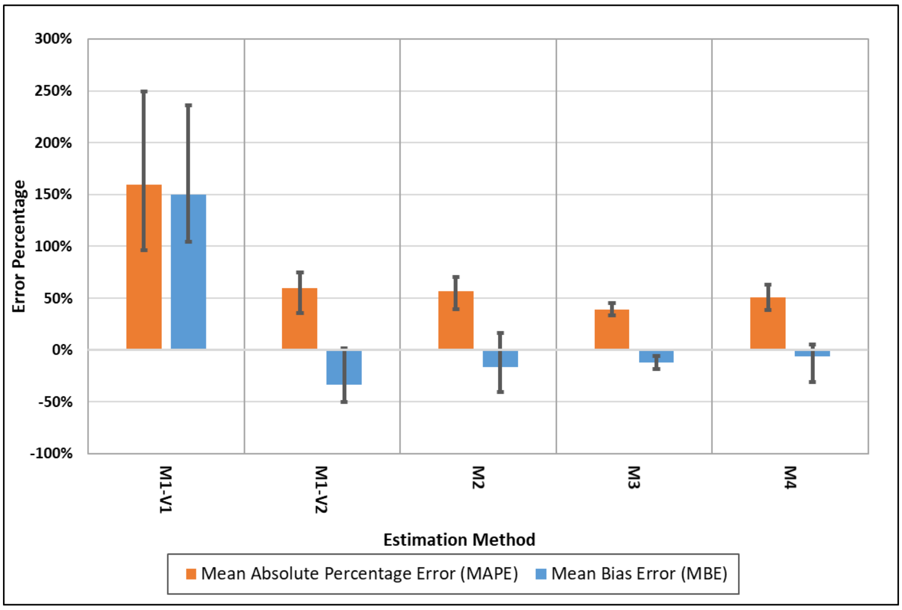

Figure 13 - Summary of each Estimation Method's error levels

Considering the outcomes of each method: the two sub-variations of M1 had contrasting results, with M1-V1 having the highest overall average error level at 159%, while M1-V2 had a considerably lower error level of 60%. The high uncertainty level of the M1-V1 method results from assuming a fixed usage level of a chiller when the actual sub-meter data shows a highly variable pattern based on a usage percentage mean of 20.8% with a variance of 252.5%. In contrast, M1-V2 uses the more variable input of the building's overall electricity usage levels for a year to first calculate the building's baseload usage. A percentage (in this case 10%) of the baseload is then deemed to be used by the DSR asset, producing a much lower average MAPE value of 60%. This result is unexpectedly low considering the method still uses a fixed proportion of the building's usage, which only considers time of day variation and results in the same half-hour prediction values being used for the entire year. The error level is still high due to this method only taking time of day variation into account, and does not consider day of year variation which will impact the estimation results of a chiller that is highly influenced by seasonality.

An average MAPE value of 56% placed M2 as the method with the second highest level of absolute error. Comparatively, however, the average MAPE is similar to the M1-V2 and M4 results. This outcome, which is based on the method outlined by Panapakidis et al. (2014), helps support usage of their profile clustering technique as the DSR estimation results are comparable to the other methods. Caution however needs to be taken with assuming this method is comparable to M1-V2 and M4 as its assumptions around the differences between profiles indicate usage of a particular electrical asset, which may be difficult to determine in different businesses.

The lowest MAPE of all the methods was M3 at 39%. The ranking of method suitability by MAPE supports research by Piette et al. (2011) where the inclusion of temperature dependency of DSR assets in predictors improves prediction. For non-weather impacted assets, other potential regression parameters could be used including time of day, occupancy levels, or operational schedules. The drawback to this method is access to historical DSR events and obtaining suitable predictor data, which could be hard to come by.

An average MAPE value of 51% placed M4 as the method with the second lowest level of absolute error. It is possible to achieve lower levels of error as demonstrated by the researchers at the Lawrence Berkeley National Laboratory - Demand Response Research Center (Dudley, 2010) who used calibrated Energy Building Models for accurate DSR forecasting. However, the calibration methods require sub-metered data for key electrical assets which, if available, could be used directly for predicting the building's DSR usage, limiting the need for using an Energy Building Model. While this method achieves comparatively good error estimation levels even without calibration, it does have the drawbacks of requiring access to detailed plans of a building and the skill and time needed to construct the model.

4.2.2 Sensitivity Analysis of Estimation Methods

The previous review of the error in the estimation methods provides a comparative analysis of methods without accounting for the uncertainty in their input values. However, the error range in DSR estimation depends not only on the estimation methodology, but also on these input uncertainties, as well as the sensitivity of method outcome to these uncertainties. Figure 14 summarises the sensitivity profiles for each method's inputs, as determined by re-running each method with adjusted inputs (see Table 46 through Table 49 in Appendix D for figure numbers). To facilitate comparison of sensitivity between methods, the charts shown in Figure 14 have been normalised. Plotting change in input variable as a percentage of the base case value against the percentage difference in estimated electricity use (MWh), Figure 14 shows varying sensitivity to inputs within and across the four methods. This section examines each method's sensitivity profiles to gain further insights into how input variation influences these.

The asset usage percentage input gradients of M1-V1 (1:1) and M1-V2 (1:1) show they are both sensitive to changes, while adjustments in the percentile value used for baseload estimation in M1-V2 have little effect (0.04:1). Altering the asset usage percentage input values for M1-V1 and M1-V2 had different impacts on the resulting MAPE outcomes across both hotels and years. The M1-V1 MAPE outcomes varied from -28.2% to 29.5%, with a consistent pattern of the MAPE value decreasing as the percentage of asset usage value lowered. This indicates that the base usage value of 50% is too high, and a lower value should be used to better represent actual usage of the chillers. The M1-V2 MAPE outcomes had a greater variance level of -11.5% to 82.8% and in contrast to M1-V1, when the asset usage percentage of the baseload value is lowered, the MAPE values increased. However, when the usage percentage is increased, the Hotel 1 MAPE values initially lower before then increasing indicating that the base value is close to optimal while Hotel 2 MAPE values continue to decrease as the usage percentage increases, indicating that a higher base value would be more appropriate. The other input for M1-V2, percentile baseload value, has a negligible effect on the MAPE outcomes with a variance range of -1.0% to 1.6% across both hotels and years, and therefore the base value of 10% is deemed appropriate.

M2 has a non-linear sensitivity profile, with each hotel and data year being impacted differently with no clear pattern. The percentage change in MAPE values resulting from the input changes has a variance range of -6% to 7% across both hotels and years. This level of MAPE variance implies that changing the number of clusters has only a small impact, and that the base value is appropriate for this application of the estimation method. The limited output variance could be the result of this method calculating the chiller usage values based on differences between cluster profiles, which means that adding or removing a single cluster will only cause the redistribution of input values into other similar clusters without causing major changes in the generated profiles.

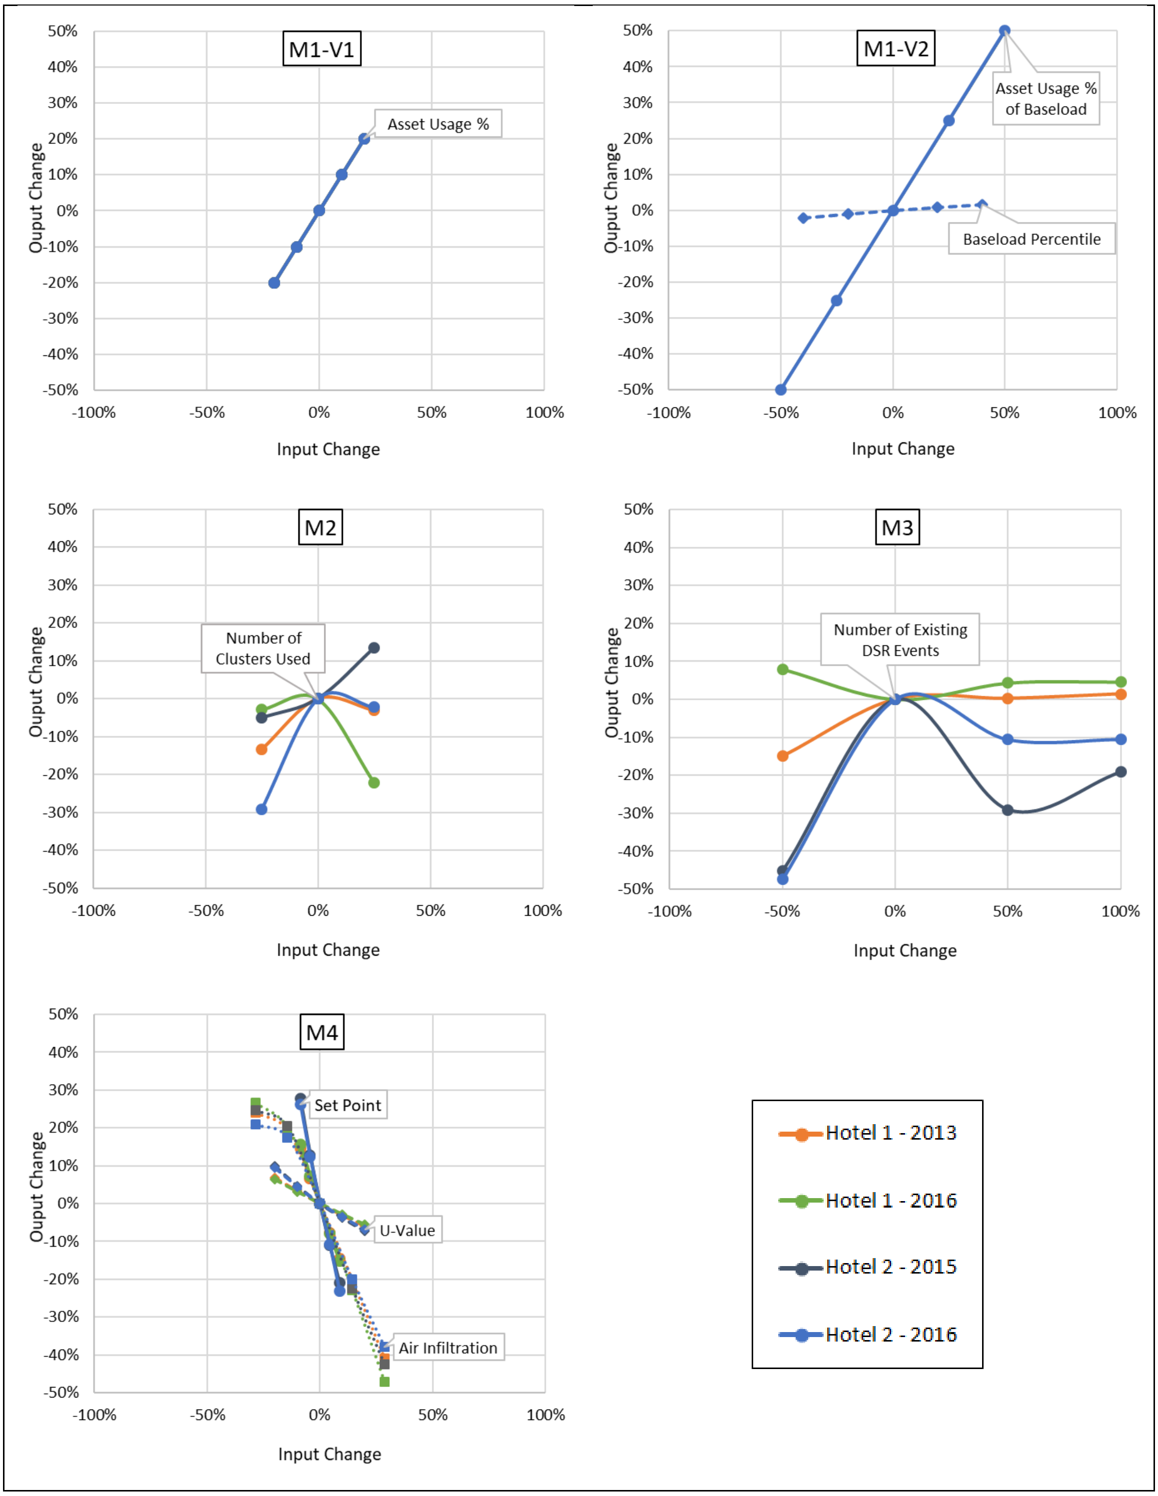

Figure 14 - Estimation Method Sensitivity Analysis Results

M3 also has a non-linear sensitivity profile that varies differently between the two hotels. The general pattern of the profile shows that when the number of historical events is lowered by 50% from 12 to 6, this has the greatest impact on estimation outputs, with MAPE values increasing by 4% and 23% for Hotel 1, and 28% and 75% for Hotel 2. When the number of events is increased to 18 and 24, the profile shows a more consistent change, except for Hotel 2 -2015. When excluding Hotel 2 -2015, the MAPE values had a minimal change range of -2% to 5%. However, Hotel 2 - 2015 showed far greater changes, with the MAPE value increasing by 45% and 28%. A potential cause of this difference could be due to the facilities manager of Hotel-2 deciding when to turn the chiller system on and off during the year. In 2015 it was turned on in April and off in October, conversely in 2016 it was turned on in May, but not turned off again. In contrast, the Hotel-1 system is left running all year with output adjusted automatically as required to meet the set point conditions. Based on the overall results of this method, reducing the number of historical events has a negative impact on the outcomes. Whereas the impact on increasing the number of events used is unclear due to the outcomes of Hotel 2 – 2015.

M4 has three different input variations of Cooling Setpoint, U-Value, and Air Infiltration. The Set Point Temperature and U-Value inputs have linear sensitivity profiles with gradients of (1:0.32) and (1:0.6) respectively, while the Air Infiltration is almost linear at (1:0.7) except for the lowest input value of 0.5 which displays as the only non-linear element on the profile. Air Infiltration changes displayed the biggest impact on output and resulting MAPE values. This is shown with the MAPE values for Air Infiltration having a variance range of -18% to 54%. In contrast, the MAPE values range for the U-Value input was -8% to 8% and the Cooling Setpoint input range was -18% to 27%. The results show how changing the Set-Point temperature and Air Infiltration rates have significant impacts on the chiller usage compared to only a minor impact from changing U-Values. This could reflect the usage of mechanical space cooling, which actively responses to temperature requirements and causes pressurised losses through Air Infiltration. The Air Infiltration input having the biggest impact does raise concern for this type of estimation method, as this is one of the hardest parameters to determine when constructing the energy building model. The other inputs can be obtained with relatively high accuracy by obtaining the Set-Point directly from the building’s current setup, and the U-Values from visual inspections of the existing construction and building plans. In contrast, the Air Infiltration rate can only be accurately obtained through a building pressure test which would be infeasible for a building of this size. Therefore, the default building model Air Infiltration rates will need to be used, and caution taken on the final outputs.

4.2.3 Cost versus Method Estimation Errors

The final set of results compares the cost of running each method against the expected level of estimation error. This comparison helps provide context to usage of the methods when balancing cost against acceptable error levels. Figure 15 maps out the links between each method's average MAPE results as per Table 45 and the estimated cost to run as per Table 22. Each method will be further examined to understand the implications of method costs and input requirements on error outcomes.

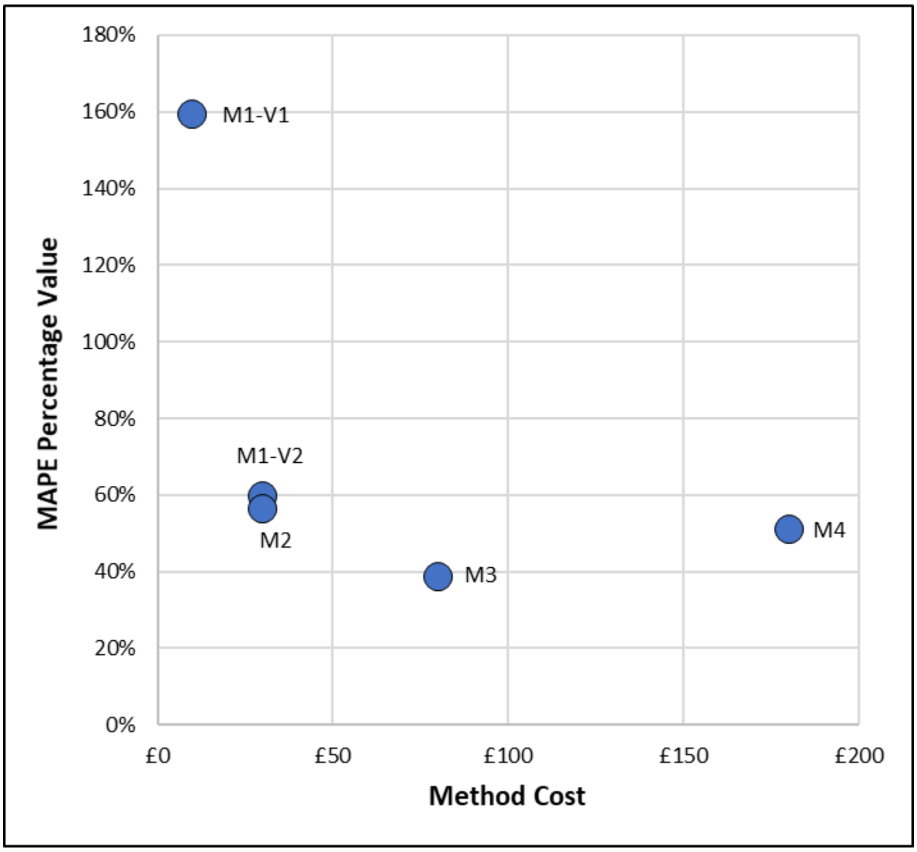

Figure 15 - Comparison of Estimation Method Error versus Cost

M1-V1 has the distinction of being the lowest cost estimation method with the highest error level. This can be directly related to the input requirement of only needing to know the asset's maximum kW rating, and then using a percentage of this for the estimation. This requires minimal time for a person to undertake, both in collecting the required information and using it to calculate the estimation. Unfortunately, the high error level implies that this method should only be used if there is insufficient time or information to undertake one of the other lower error proceeding methods. In comparison, M1-V2 reduces the error level by two-thirds compared to M1-V1 while costing 3 times more to run. While M1-V2 is more expensive than M1-V1, it is still comparatively cheap compared to all the methods tested. This method also uses relatively accessible data of the building's electricity usage records, which in the UK is available in a half-hourly format for any business with peak electricity usage of 100kW or greater.

M2 is the third equally cheapest method to run due to the primary input requirement being the building's half-hourly electricity usage records. It also has the fourth lowest error level, which makes it a highly recommended method if both cost and accuracy are taken into account. However, as discussed previously, this method's usage of clustering means that care needs to be taken on its application to suitable buildings and assets.

M3 achieved the lowest error level of all methods tested at 39%. However, it also has the second highest cost at £80, which is a result of requiring two expensive input requirements. Firstly, it uses detailed historical air temperature readings over a year for the building's location, which requires paying for access to the necessary weather archive. Secondly, it uses previous DSR event outcomes which require time to obtain from the building users, and then formatting and verifying before using. The requirement for existing DSR event outcomes is also a key limitation of this method. The ability to provide this information will be dependent on whether the site has previously been involved in DSR and is able to provide information on the kW reductions occurring during the events.

M4 had the highest cost at £180 with the second lowest error level of 51%. The high cost is primarily due to the time required to model the building in the building energy modelling tool. As the resulting error level is similar to M1-V1, M2 and M3 methods, which are significantly cheaper to run, this method is not recommended. Although a potential justification for using this method would be if multiple assets within one building were being estimated, thereby reducing the individual assessment costs while providing a combined view of the building's potential.

4.3 Chapter Conclusion

This chapter has undertaken an examination and comparison of four non-domestic DSR estimation methods to provide insights into uncertainty levels based on the input requirements. The examination was performed by using each method to estimate the DSR potential of HVAC chiller assets at two hotels over two years. The estimation outcomes were then compared against the chiller's actual sub-metered usage records by calculating MAPE and MBE values to understand each method's level of estimation error. The results showed a wide range of estimation errors. Method 1 - Sub-variation 1 yields the highest error level MAPE of 159%, while the lowest error level MAPE of 39% was achieved with method 3. While method 3 could be a recommended approach based on its low error level alone, it is unfortunately not that simple due to information input considerations. Based on this chapter's findings, each method requires review to understand the implications of input requirements on outcome uncertainty. These findings can be summarised as follows:

- **Method 1** sub-variation 1 has the lowest informational requirement and cost of £10 to use based on only needing to know the maximum kW rating of the asset being assessed to apply this method. However, the drawback of this low informational requirement is the highest error level of all methods at 159%. Sub-variation 2 achieved a much lower error level of 60% by using the building's half-hourly electricity records that increases the usage costs to £30. The sensitivity results for this method showed a high impact on the outcomes based on variations of the inputs. This means that the error results might differ substantially when used in other scenarios. Therefore, the error levels reported in this research for method 1 need to be used with care when deciding on suitable assessment approaches.

- **Method 2** had the second highest error level of 56% while being the third cheapest to run at £30 through clustering of the building's half-hourly electricity usage data. The sensitivity analysis of this method showed a medium to low impact on error levels arising from changes in the primary input of how many clusters are used. These results indicate that baseline comparison is a suitable method for assessment, though it has two limitations that need to be fully understood by users to ensure valid results. Firstly, it requires the user to select the appropriate number of clusters, which is open to individual interpretation. Secondly, this method will only work on electrical assets that have enough variation within the building's overall usage to be identified by the clustering.

- **Method 3** had the lowest overall error level of 39% with the second highest cost of £80. The low error level makes its utilisation of historical DSR event outcomes an attractive method. However, its practical usage is limited as it requires the building to have previously undertaken DSR and have access to historical DSR events outcomes. The sensitivity analysis also showed a significant increase in error if less than 12 historical event records over a year are available for analysis. In new DSR markets these limitations may restrict usage of this method. Even in established markets it could be difficult or time-consuming to obtain any adequate historical information from the existing DSR aggregator.

- **Method 4** had the second lowest error level at 51% but had the highest cost of £180, which is over twice that of method 3, the next most expensive, as a consequence of the amount of time required to develop a building energy model. While this method had the second lowest error level, it is only slightly lower than many other cheaper options and method 2, for example, costs 6 times less with only a slightly higher error level of 56%. The usage requirements of this method also restrict its practical application given its reliance on detailed building plans and the skills to develop building models. The importance of having the right information and skills is highlighted by the sensitivity analysis, which showed major impacts from variations in temperature set-points and air infiltration model values.

These findings have three key implications on selection of DSR estimation methods. Firstly, the wide range of error levels means the outputs of these methods will need to be carefully considered when being used to make decisions about the suitability of buildings for DSR. Secondly, care needs to be taken in ensuring accurate input selection as sensitivity analysis demonstrates that adjusting the inputs on most methods will result in large variations to the outputs. Thirdly, this research tested four methods using HVAC chillers in hotels only. Therefore, other assets and businesses may result in different error outcomes and caution needs to be taken before this research is used to select estimation methods outside of this scope. This final implication highlights a potential future area for research which would entail re-running the method comparisons on different DSR assets and businesses to understand the different impacts on estimation outcomes. The impacts of these findings are further reflected on in the chapter 6 thesis conclusion.

5 Development of a Profile Based DSR Estimation Method

This chapter covers the final research objective '*To develop and evaluate a model that uses asset usage profiles to reduce the uncertainty of DSR potential estimation during an aggregator's assessment process*'. Chapter 4 identified four existing DSR estimation methods and then, through a comparison, showed that these methods have a range of output errors, from a MAPE of 39% to 159%. The comparison also showed variations in the cost of running the methods, ranging from £10 to £180, with the lowest MAPE method costing £80. The existing estimation methods are also revealed as deterministic in nature, and therefore unable to provide any measure of certainty in their outputs. Accordingly, obtaining any accuracy in understanding the nature and scale of the output uncertainty for existing methods would require retroactively evaluating estimation results against actual data once the new site has gone live. Yet this is inefficient as a solution for improving the viability and market penetration of DSR solutions.

Therefore, as an alternative, this chapter develops a new estimation method that aims to help address existing estimation methods' limitations. The Profile DSR Estimation Method, or '***profile method***', established in this chapter uses detailed usage information (from sub-meters for example) of electrical assets to create a set of load profiles that represent common usage patterns. These profiles can then be used to determine the likely usage levels for similar assets at sites where detailed information is not available. Profiles were selected for use to create a new DSR estimation method because a review in section 2.4 of building energy usage estimation approaches demonstrated load profiles as having the greatest opportunity for development as a new DSR estimation method.

Once the DSR estimation profiles have been created, this method aims to provide as benefits: (i) only requiring very basic information to perform a DSR potential assessment for a new site (namely the site's business category and maximum kW ratings for potential DSR assets), (ii) requiring minimal effort to apply, as the usage profiles are applied automatically once the basic information is entered, and (iii) providing the profile method's user with the ability to manage and understand the estimation uncertainty. The primary limitations of the profile method include: (i) creating reliable usage profiles will first require, for example, hard to obtain sub-metered data for potential DSR assets, which can then only be used for similar assets and business categories, and (ii) developing the usage profiles and a software application that others can use for applying them in estimating a site's DSR potential will require time and specialised skills, and therefore investment.

By seeking to develop a new estimation method through creating and then applying load profiles, this chapter builds upon existing literature's published approaches towards analysing electricity demand, which also involves using load profiles. At a macro level, for example, load profiles are already used by the electricity industry to understand country-level demand for electricity. Elexon (the UK's electricity settlement service) has already created for the UK market eight expected usage profile classes for domestic and non-domestic sites that lack half-hourly metering, in order to enable demand forecasting for these sites based on a temperature based regression model (Elexon, 2013). These Elexon profiles are generated by using data captured from up to 2,500 half-hourly meters which are installed across sites throughout the UK. The eight expected usage profiles are then created by averaging from this data the anticipated demand for each profile group before using regression analysis to determine temperature driven coefficients.

In contrast, Räsänen et al. (2010) state that site usage profile classification methods like Elexon's, which rely on site characteristics and annual electricity usage, increase uncertainty in demand forecasting due to risks of incorrect profile assignments or changes in usage patterns. However, these authors still make use of load profiles, but claim that these should be assigned based on data-driven principles whereby recent improvements in usage monitoring can enable the actual usage of a site to be applied when assigning load profiles. Räsänen et al. demonstrate this theory by creating profiles using clustering methods that identified 19 groups from 3989 sites in Finland, and then using new and existing profiles to forecast electricity usage at 230 sites that were not used in the clustering. The authors deemed their results to show major improvements in the accuracy of estimation due to better use of load profiles, as the index of agreement mean increased from 0.478 to 0.627.

Load profiles have also previously been adopted in enabling estimation at a more granular scale, specifically in assessing electricity demand at room and end-usage levels. Liddiard (2014) used energy end-use survey information from over 300 non-domestic buildings to produce room-scale energy usage profiles for 16 room types. These profiles were then applied using premise type and room floor space parameters in order to calculate overall energy usage. Widén et al. (2009) used time-use data to generate load profiles for domestic electricity and hot water usage. These researchers were able to create load profiles for different end-use categories (e.g. computer usage, cooking, watching TV) by combining the electricity actually used by standard appliances with the time-use activities. These load profiles were applied to households to determine overall electricity and hot water usage, and these results were then compared against a set of households that had detailed measured usage data. Widén et al. found that these profiles provided realistic reproductions for overall electricity demand, yet less accuracy for hot water usage due to difficulties with the time-use data, as it seldom matched the actual time that hot water was in fact used by households in appliances like dishwaters and washing machines.

While the published research suggests that the usage of load profiles is extensive in the energy industry, there is no known literature on using load profiles to assess a site's DSR potential as described in this thesis. The closest relevant research arises in the context of country-level DSR predictions. This is illustrated in a report by Element Energy for OFGEM (2012), which reviewed the non-domestic DSR potential for the UK. Element Energy determined their conclusions about the UK's overall DSR potential by first creating industry level profiles (for offices, retail etc), which then split out the overall end usage values for the areas of catering, computing, heating, hot water, HVAC, and lighting. Yet these profiles were limited, as the end-usage levels were calculated using a single percentage value for each area and sector only, in turn resulting in an assumed constant rate of usage across the day then only adjusted based on site level half-hourly electricity usage data for the sector and seasonal adjustments. The report's authors note this as a limitation resulting from the lack of available sub-metered data for the end-usage areas. They also noted that they could not create sector-specific profiles for hotel and catering or communications and transport due to a lack of available site level half-hourly electricity data.

Against this context, for the purposes of this thesis, previously published research into load profiling for the energy industry will be used to help develop a new profile-based method, which seeks to improve the reliability of DSR estimation. The development of the profile method as a new approach for estimating a site's DSR potential is undertaken using a seven-stage process, outlined in Figure 16. While this chapter addresses each stage of the development process in detail, by way of a high-level overview: the development process starts with obtaining the primary information input of sub-metered usage data for the same types of electrical assets as used in selected categories of businesses (for example, HVAC chiller assets at hotels). This data is then normalised by calculating the asset's usage as a percentage of its maximum potential. The normalised sub-metered data is then clustered into comparable groups of daily usage patterns. Each group is then formed into a profile that represents the upper, median, and lower usage boundaries of the group's daily patterns. The profiles are then assigned a calendar based predictor that best represents how the data was clustered (for example, Week-of-Year, Month-of-Year, etc).

The profiles are then utilised by applying these to similar assets and business categories to estimate their potential for DSR. For this development the Python 2.7 programming language was the primary tool used. Python was selected for programming based on its versatility, performance, support, capability for large dataset handling and range of supported modules. Using Python for this type of analysis reflects a 2017 survey by KDnuggets, which showed that Python had become the leading tool of choice with 52.6% of 2,900 surveyed members of the analytics and data science community using it (KDnuggets, 2017).

This chapter is divided into nine sections. Sections 5.1 to 5.7 each address one of the seven development stages outlined in Figure 16. Section 5.8 reviews and discusses the outcomes of this method's development, and then compares its results to the outcomes from the existing estimation methods that were reviewed in Chapter 4. Section 5.9 concludes this chapter's findings.

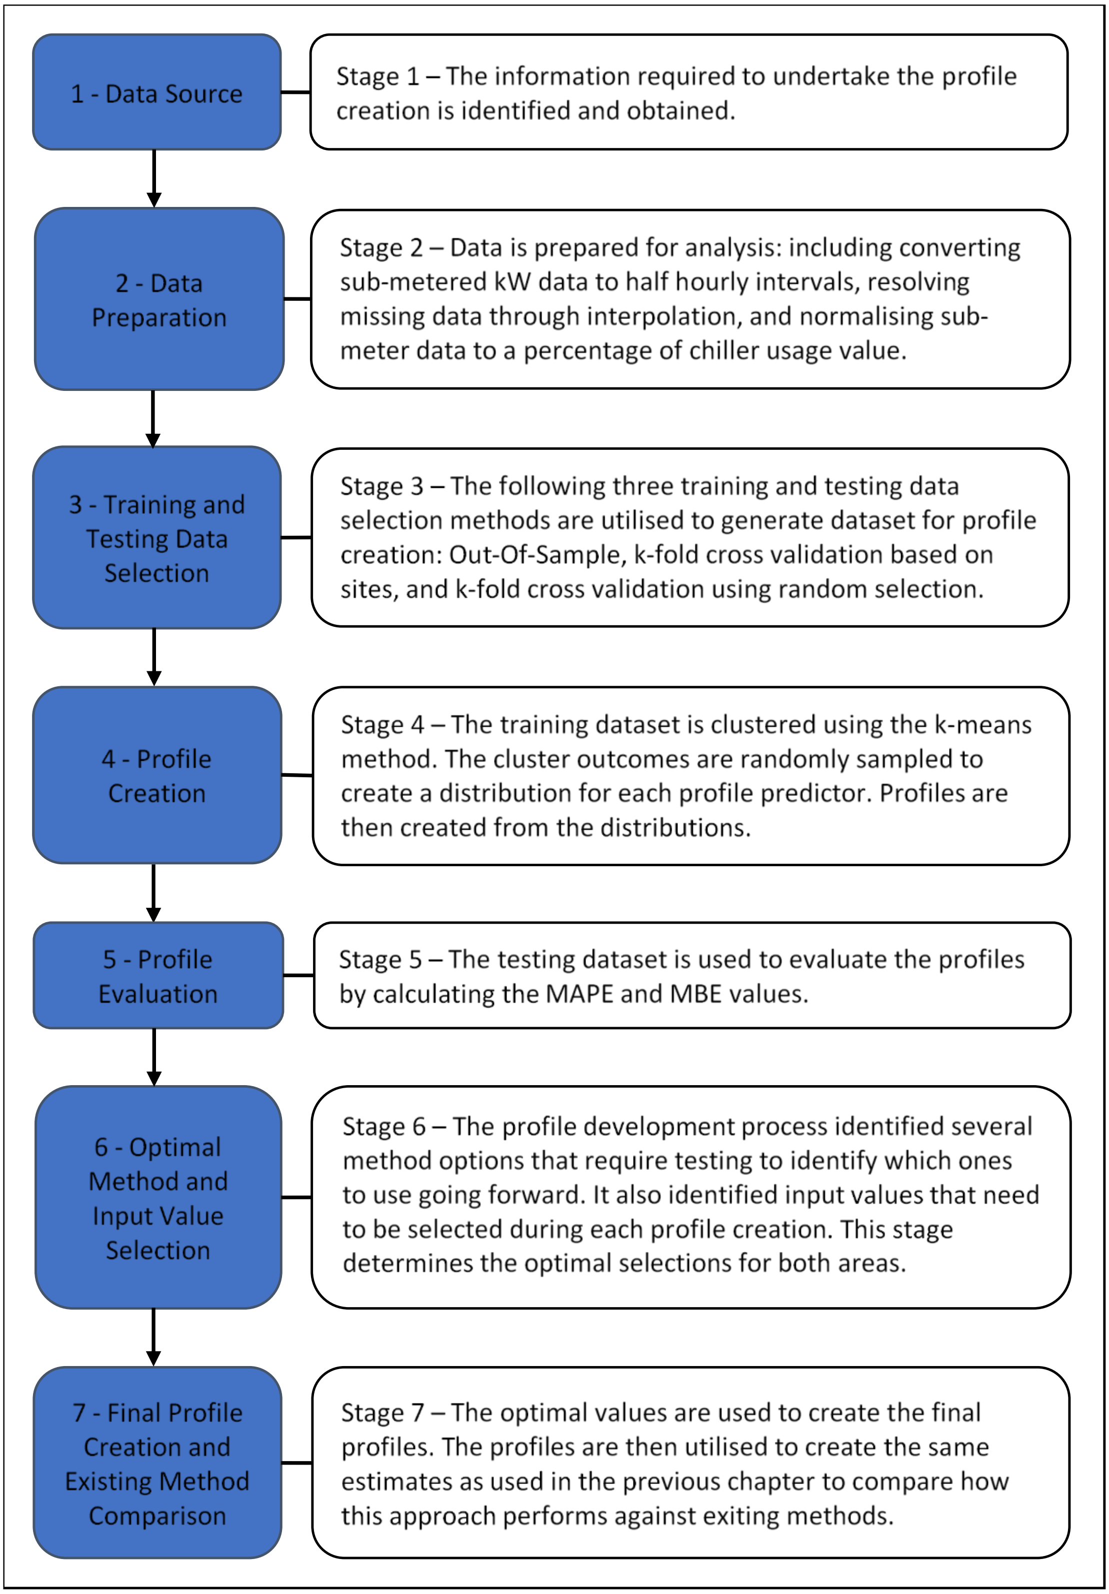

Figure 16 - Summary of Development Stages for the New Profile Estimation Method

5.1 Stage 1 - Data Sources

The first stage of developing the profile method requires determining and accessing the raw data required to generate the asset usage profiles. The data was obtained from the research partner, KiWi Power. A review of KiWi Power's available data identified that HVAC chiller units in hotels provided the best dataset for profile development as sub-metered usage records that started in 2012 were available for use and section 2.1.3 identified them as being suitable for DSR due to their operational flexibility. As a consequence of KiWi Power working with a hotel chain to assess the potential of DSR, the available data spanned multiple years across five hotels. As the hotel chain was enabled for the STOR Turndown DSR product, sub-metering information was required and collected for each site to ensure accurate measurement of the assets' performance during DSR events to meet the National Grid STOR programme's metering requirements, which require minute interval kW readings (National Grid, 2017f). The reading is taken using the KiWi Power PiP hardware (see section 3.1.1 for additional information about the PiP) by measuring the electric current ampere (using a current transformer) and voltage that are then multiplied to calculate the real power value. The minute interval kW value is then obtained by averaging the real power usage of the HVAC chiller over a minute and rounding to the closet kW. The PiP hardware has been certified for DSR metering by the National Grid, which means that the National Grid deems the hardware's measurements as being within an acceptable level of error.

The main HVAC chiller units comprise the primary assets used for DSR in the hotels supported by KiWi Power. As outlined in section 2.1.3, HVAC chillers are deemed suitable for Turndown DSR as they can be turned off or down for up to an hour without occupants noticing, and offer a large centralised electricity using asset that is straightforward to enable for DSR. These assets are straightforward to enable as HVAC chillers normally have a specific on/off DSR switch that effectively turns off the system while activated. The switch can either be manually operated or controlled via an external system. For the chiller systems used in this research, KiWi Power used the same PiP device to monitor usage and provide remote activation of the DSR function.

The KiWi Power web-based management system 'KiWi Operation Management Platform' (KOMP) was used to acquire information about the hotels and the sub-metered usage data of their HVAC chillers. Table 24 summarises the information gained from the KOMP system about the chillers in each hotel. The 'Meter ID' information was used to extract the actual minute interval kW usage data for each chiller into a CSV file. The 'Chiller kW Rating' value will be used later in this stage to normalise the kW values to a percentage of chiller usage. The 'Days Containing kW Usage Data' indicates the number of days that usage records of the chiller are available for in the KOMP system. The difference in the number of recorded days of data between hotels resulted from the hotels monitoring being temporarily stopped after an initial trail operation throughout 2013. Monitoring was reactivated again once each hotel signed a commercial STOR DSR contract. This resulted in mixed levels of usage records for each hotel, as shown in the different number of available data days for each hotel. Table 28 illustrates which months the data was recorded in.

Table 24 - Details of Sub-Metered Hotel Chiller Used for Profile Creation.

Column information:

Meter Input ID - the meter's unique identifier
Description - details of the chiller
Chiller kW Rating - The chiller's maximum potential kW usage amount
Days Containing kW Usage Data - The number of days with recorded usage data.

Meter ID	Hotel Name	Description	Chiller kW Rating	Days Containing Usage Data
MA_RYL_C1	Bristol Royal	Chiller 1	111	1115
MA_RYL_C2	Bristol Royal	Chiller 2	111	1115
MA_RYL_C3	Bristol Royal	Chiller 3	111	1115
MA_KIN_C1	Kensington	Main Chiller	132	952
MA_MDV_C1	Maida Vale	Main Chiller	135	365
MA_REG_C1	Regents Park	Main Chiller	290	973
MA_WRS_C1	Worsley Park	Main Chiller	86	607

5.2 Stage 2 - Data Preparation

The second development stage covers preparation of the chiller usage data that will be used for analysis and creation of the usage profiles. This preparation includes understanding the level of missing data values, converting the minute interval data into half-hour intervals, and finally normalising the usage data.

5.2.1 Understanding and Managing Missing Values

An understanding of the level of missing values needs to be gained before the data acquired from KOMP can be prepared for analysis, to decide on appropriate actions to address any gaps. To determine the level of missing values, the minute interval kW readings datasets as extracted from KiWi Power's KOMP system were analysed to identify any gaps in the readings. The results of this analysis are shown in Table 25, in the 'Number of Missing kW Minute Readings' column. The results show that overall 0.16% of the minute kW readings are missing. Part of the data preparation requires converting the minute kW readings into half-hourly average kW values (as per the next section 5.2.2). Therefore, to understand the impact of this conversion on missing values, the data was analysed to determine how many half-hour periods contained at least one missing value. The results of this analysis are shown in Table 25, in the 'Number of Half-Hour Periods with Missing Readings' column. The percentage of missing half-hour periods is 0.19%, 0.03% higher than the minute level missing readings due to all readings in a half-hour period being excluded even if there is only one missing kW minute reading.

Table 25 - Overview of Missing kW Readings

Hotel	Days with Readings	Number of kW Minute Readings	Number of Missing kW Minute Readings	Percentage of Missing kW Minute Readings	Half-Hour Periods Without Missing Readings	Half-Hour Periods with Missing Readings	Percentage of Periods with Missing Readings
Bristol Royal	1,115	1,603,192	2,408	0.15%	53,424	96	0.18%
Kensington	952	1,368,601	2,279	0.17%	45,610	86	0.19%
Maida Vale	365	525,600	0	0.00%	17,520	0	0.00%
Regents Park	973	1,398,713	2,407	0.17%	46,609	95	0.20%
Worsley Park	607	871,753	2,327	0.27%	29,045	91	0.31%
Overall	**4,012**	**5,767,859**	**9,421**	**0.16%**	**192,208**	**368**	**0.19%**

As the later stages of the profile development rely on daily usage patterns, any missing values will invalidate that day's data. Therefore, the number of half-hourly periods with missing readings were counted by date to determine how many days were impacted. Considering the days that had missing values, a consistent pattern was identified whereby the gaps all occurred on the same days and times across the hotels, as seen in Table 26. To understand the reasons behind these missing readings, the results in Table 26 were discussed with the KiWi Power operations and technical teams. Their feedback identified three main reasons for missing values: (1) the KOMP system crashed and required restarting, (2) the database reached its maximum size and required clearing before more data could be saved, (3) general communications issues resulting in readings being lost.

Table 26 - Number of Half-Hourly Periods Missing Readings per Day per Hotel
(Note – 'Offline' means the Hotel was not being used for DSR on these dates and therefore no kW readings recorded)

Date	Bristol Royal	Kensington	Maida Vale	Regents Park	Worsley Park
26/01/2016	1	Offline	Offline	1	1
03/02/2016	3	Offline	Offline	3	Offline
20/03/2016	1	Offline	Offline	1	Offline
14/06/2016	1	Offline	Offline	1	1
20/07/2016	1	Offline	Offline	1	1
22/07/2016	1	Offline	Offline	1	1
26/07/2016	1	Offline	Offline	1	1
06/09/2016	3	3	Offline	3	3
07/09/2016	25	24	Offline	24	24
19/11/2016	1	1	Offline	1	1
07/12/2016	1	1	Offline	1	1
08/12/2016	6	6	Offline	6	6
23/12/2016	1	1	Offline	1	1
27/12/2016	2	2	Offline	2	2
28/12/2016	2	2	Offline	2	2
24/02/2017	6	6	Offline	6	6
04/04/2017	17	17	Offline	17	17
16/06/2017	23	23	Offline	23	23
Total	**96**	**86**	**0**	**95**	**91**

While the dataset contains limited missing values that could be removed without materially impacting the analysis, it is important to understand how missing values could be addressed, as different datasets might contain higher levels of missing values than those presenting within the dataset observed for the purposes of this thesis. Managing missing data in time series is well researched, expressly in the area of climate data, and this research can be drawn on to understand potential options (Sluiter, 2009). For the purpose of this research, it was determined that linear interpolation would be utilised as it has been extensively used as a standard approach for dealing with time series missing data, and offers the benefits of simple and understandable calculations that can be used for varying dataset sizes with no additional data requirements (Meijering, 2002).

This method works by effectively creating a straight line between two known data points in a series, and then using this line to calculate any missing values that exist between these two points. Its main limitation is that it can be inaccurate when the missing data points would not have fitted on a linear line, as is the case, for example, where a data profile contains peaks and troughs throughout the measurement period. If the two reference points sit either side of the peak or trough, then the resulting linear interpolation line will result in the peak or trough being removed.

To understand the impact of using linear interpolation on the acquired half-hourly chiller usage dataset, an assessment was undertaken to determine the level of error that would be experienced over different scales of missing values. This analysis was performed applying the following steps:

1. Using Python, a hotel chiller dataset containing one year of non-missing half-hourly kW usage values is extracted from the available kW records.

2. The dataset was iterated through sequentially, with one value in the dataset being temporarily removed at a time. For each removed value the linear interpolation method was used to estimate the missing point. This was done by calculating a straight line between the known values before and after the missing point. Using the slope of the line an estimated value of the missing point is calculated. After the missing point has been estimated using interpolation the real value is then returned to the dataset before the next sequential value is removed.

3. To understand the level of error introduced using linear interpolation, the MAPE value of each estimated point is calculated. The MAPE value (as described in the previous chapter 4.1.3.1) first calculates the difference between the estimated and actual half-hour values and then divides the result by the actual value. The absolute value of the resulting outcomes for all estimated points is summed and divided by the number of estimated points. The resulting value provides a percentage indication of the level of fit between the estimated and actual values, with lower being better.

4. Repeat steps 1 – 3 with the number of sequential missing values increasing by 1 until a maximum of 48 missing values (i.e. an entire day) have been tested. The MAPE results for each sequential missing value length are shown in Figure 17.

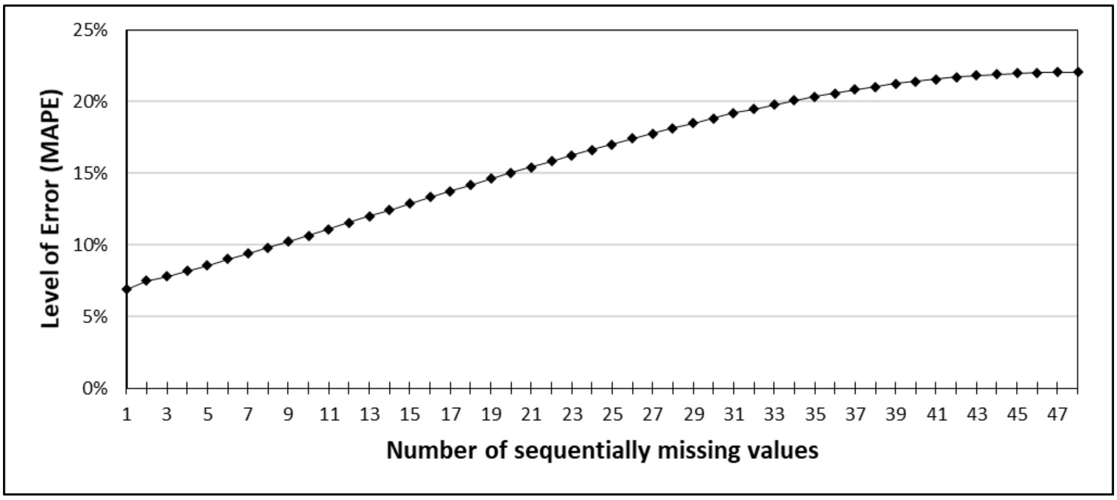

Figure 17 - Error levels resulting from using interpolation across a range of missing value lengths from a dataset containing 365 days of data.

The linear interpolation assessment results, as set out in Figure 17, show a gradual increase in error level as the number of sequentially missing values increase. The error level gradient starts to decrease once the number of missing values reaches 33, and then demonstrates almost no change at 44 and beyond. The lowest error value of 6.7% resulted from missing one value which, based on Table 27, occurred the most often in the dataset. The highest number of sequentially missing values at 25 occurred once and, based on Figure 17, would result in having a potential 17% error level if fixed using linear interpolation.

Table 27 - Count of Sequentially Missing Half-Hourly Values

Sequentially Missing Half-Hourly Values	Count of Occurrences
1	29
2	8
3	6
6	8
17	4
23	4
24	3
25	1

Based on these results, three options are available for managing the missing data used in this research. Option one is to exclude all days that have missing values. Based on Table 26, if all days that contain at least one missing reading are excluded, then 63 of the 4,012 days of data (1.57%) would be removed. Option two would be to use linear interpolation to estimate the missing readings, which would result in introducing error levels on those specific values of between 6.7% to 17%. Option three is to use a mix of option one and two, whereby some days are fixed by using a set level of missing values determined using linear interpolation (e.g. fix any day that has only 1 missing value), with all other days with missing data being excluded. However, it is unclear how each of these options will impact the final profile error levels. Therefore, as the purpose of developing the profile estimation method is to reduce uncertainty in DSR estimation for new sites, it was determined that for the purposes of this research, each option would be assessed during the optimisation stage 6 (section 5.6) to determine which option is most suitable for creating DSR estimation profiles.

5.2.2 Converting Sub-Metered Data to Half-Hour Intervals

The extracted chiller kW usage data was recorded in minute intervals as previously described in section 5.1. The decision to convert the data into half-hourly usage intervals was due to: (1) the UK electricity system works on half-hourly settlement periods, and half-hourly usage intervals therefore comprise standard units of measurement recognised in the industry (Elexon, 2014), (2) half-hourly usage intervals enable easy comparison of estimation outcomes against the site's overall actual half-hourly electricity usage levels, and (3) this approach uses the same data sample size as used in Chapter 4's comparison of existing DSR estimation methods. To undertake this conversion, the minute kW values for each half-hour were summed and divided by 30. If there were less than 30 values present, then that half-hour period was marked as having missing values and excluded. This resulted in 368 half-hour periods being marked missing, as per Table 26.

5.2.3 Normalising the Sub-Meter Data

Once the kW usage data has been converted to half-hourly periods, it was then normalised as a percentage of overall chiller capacity. As the chiller at each hotel is a different size, as shown in Table 24, the kW usage information cannot be directly compared and used for profiling without normalisation. The data is normalised by calculating the asset's percentage level based on its maximum kW rating. By converting the kW usage data into percentage usage values, chillers of different sizes can be directly compared.

Before the data could be normalised, the Bristol Royal Hotel dataset required additional adjustment as this site has three identical chillers feeding the same HVAC distribution system. As the chillers worked in parallel, their datasets could be combined to form one virtual chiller with a maximum potential rating of 333kW based on the sum of each individual unit, as per Table 24. The normalisation of kW usage values to percentage of chiller usage was then undertaken by dividing each half-hourly kW usage value by the applicable chiller's maximum kW rating.

5.3 Stage 3 - Training and Testing Data Selection

The third stage outlines options for splitting the normalised half-hourly chiller percentage usage data into training and testing datasets that will then be used in stage four for creating the profiles, and then in stage five to evaluate the profiles. The reasons for selecting three selection methods are described before outlining how these are implemented.

5.3.1 Review of Methods for Selecting Training and Testing Data

Selecting how the data will be split into training and testing datasets is an important consideration in ensuring that evaluation of the proposed method is valid. When reviewing potential options, the research by Bergmeir, Hyndman, & Koo (2015) on cross-validation for evaluating time series prediction notes that '*when it comes to time series prediction, practitioners are often unsure of the best way to evaluate their models*'. The authors attribute that this is due to time series data serial correlation, possible non-stationarities, and general practitioners consider that future data should not be used to predict the past (i.e. using new data in a time series to formulate a prediction of older data in the same series). Due to the apparent uncertainty about how best to evaluate models, these authors found that most people defaulted to using the out-of-sample (OOS) approach. This reflects previous research that also found the OOS method as the traditional default approach for time-series evaluation (Christoph & José, 2012). However, both research groups proposed and proved that cross-validation methods can provide better results than OOS if the information is contained in blocks of data, instead of splitting each individual record in order to maintain the continuous nature of time series.

As the literature shows differences of opinions regarding which method to use for training and testing dataset selection, for this research, three methods that cover the traditional and novel cross-validation approaches are reviewed. The first method is the OOS approach that Bergmeir, Hyndman, & Koo (2015) deem to be the traditional method used for evaluating time series data. It works by selecting the testing dataset as a percentage of the most recently recorded values. The remaining older values in the time series are then used as the training dataset.

Table 28 illustrates this approach, using the latest 25% of records for each hotel for the testing dataset and the remaining 75% for a training dataset. The percentage splits to use for OOS will vary depending on the method's user, yet a general range of between 20-30% has been deemed as an acceptable amount for the testing dataset, and therefore for this research, the mid-point of 25% was selected (Nau, 2017; Research Gate, 2017; Stack Overflow, 2017).

The second and third methods 'K-fold Cross Validation Using Sites' and 'K-fold Cross Validation Using Random Selection' both use cross-validation for creating the training and testing datasets. Cross-validation works by first dividing the dataset into k subsets based on a predefining attribute (Christoph & José, 2012). A testing dataset is then formed of one k subset with all other subsets then being used for the training dataset. The process is repeated until each subset has been used as the training dataset. The results of the evaluation using each combination of training and testing datasets are then averaged to provide the overall outcome. The 'K-fold Cross Validation Using Sites' method applies this approach by defining each site (in this case each hotel) as a k subset, which results in five sets of data for training and testing. Table 28 illustrates this method by defining the Bristol Royal Hotel as the first k subset to be used for the testing dataset, with the remaining four hotels being used for the training dataset. The 'K-fold Cross Validation Using Random Selection' method applies this approach by randomly assigning each day in the dataset with a value between 1 to 4 based on creating four k subsets, which results in four sets of data for training and testing. Table 28 demonstrates this approach by assigning the random values by month (for the purposes of the example) and then, for the first data set, selecting all months assigned to group 1 as the testing dataset, with the remaining groups 2-3 then being used for the training dataset.

Table 28 - Example Illustration of Training and Testing Data Split Methods

The example shows creation of training datasets in blue and testing datasets in yellow for each method. Method 3 demonstrates random assignment of data to groups 1-4 with group 1 being used as the testing dataset.

		Selection Method 1: Out-Of-Sample using last 25%					Selection Method 2: K-fold Cross-Validation 1 site at a time					Selection Method 3: K-fold Cross-Validation 4-way random split				
Year	Month	Bristol Royal	Regent Park	Worsley Park	Kensington	Maida Vale	Bristol Royal	Regent Park	Worsley Park	Kensington	Maida Vale	Bristol Royal	Regent Park	Worsley Park	Kensington	Maida Vale
2013	1											3			1	1
	2											3			1	4
	3											4			3	1
	4											3			3	2
	5											4			3	1
	6											2			3	2
	7											4			1	2
	8											2			1	4
	9											1			4	4
	10											1			1	4
	11											1			1	1
	12											1			3	3
2014	1															
	2															
	3															
	4															
	5															
	6															
	7															
	8															
	9															
	10															
	11												1			
	12												2			
2015	1												1			
	2												1			
	3												1		4	
	4												2		2	
	5												1		4	
	6											1	2		4	
	7											3	4		1	
	8											2	3	1	2	
	9											1	1	3		
	10											1	2	2		
	11											4	4	3		
	12											2	1	3		
2016	1											1	2	1		
	2											1	4			
	3											1	3			
	4											2	1			
	5											2	1	2	4	
	6											1	3	2	3	
	7											3	1	3	2	
	8											3	2	1	4	
	9											2	3	1	3	
	10											2	4	3	3	
	11											4	4	2	3	
	12											3	2	1	3	
2017	1											3	3	2	3	
	2											2	1	1	3	
	3											4	4	1	4	
	4											2	3	3	4	
	5											1	2	2	1	
	6											2	3	3	2	

The research by Bergmeir et al. (2015) and Christoph & José (2012) shows that each of the three described methods have different benefits and limitations, as summarised in Table 29. The first main difference between the OOS and cross-validation methods is that the former only uses one part of the dataset for testing, which has the benefit of ensuring independence between the data used for training and testing, yet means that it is susceptible to atypical data. The cross-validation methods loses independence as the whole dataset is used to train and test the dataset, yet Bergmeir et al. (2015) have proved that it can provide greater accuracy if the data is kept in logical blocks of common data (for example as daily usage blocks) to maintain the continuous nature of time series data. As these methods represent the traditional and novel approaches for managing time series data, for the purposes of this research it was decided that all approaches would initially be used during the profile method's development process to create training and testing datasets for profile creation. This then enables the outcomes of each method to be critically reviewed during stage 6's (section 5.6) optimisation process to finally determine which method to recommend as optimal for profile creation.

Table 29 - Evaluation Methods Benefits and Limitations Matrix

Benefits and limitation descriptions based on research by Bergmeir et al. (2015) and Christoph & José (2012):

B1 – Traditional accepted approach
B2 – Simply data selection method
B3 – Training and testing data remain independent
B4 – Multiply evaluations using the whole dataset
L1 – Susceptible to atypical usage patterns within the dataset
L2 – Does not utilise the full dataset for training
L3 – Single evaluation based on one part of the dataset
L4 – Loss of independence between training and testing data
L5 – Non-traditional time series approach with limited research

	Benefits				Limitations				
Method	**B1**	**B2**	**B3**	**B4**	**L1**	**L2**	**L3**	**L4**	**L5**
Out-Of-Sample	●	●	●		●	●	●		
K-fold Cross Validation Using Sites		●		●	●			●	●
K-fold Cross Validation Using Random Selection				●				●	●

5.3.2 Creating Training and Testing Datasets

The three data selection methods (OOS, K-fold Cross Validation Using Sites, and K-fold Cross Validation Using Random Selection) outlined in the previous section 5.3.1 are each used to create training and testing datasets from the normalised half-hourly chiller usage data. The training and test datasets are then each used during the Profile Creation (section 5.4) and Evaluation (section 5.5) stages. The steps undertaken to generate the training and testing datasets for each method comprise:

1. A master dataset of daily half-hourly chiller percentage usage values of all the hotels is created based on the data generated in stage 2 (data preparation), as described in section 5.2. The dataset is constructed so that each row represents the daily usage values for a chiller (e.g. each row contains 48 half-hourly usage values, date and hotel).

2. The Method 1 OOS training and testing datasets are created by marking the first 75% of daily usage records from oldest to newest in the dataset with the number 1, then marking the remaining 25% with the number 2. The dataset is then split into a training dataset that contains all values marked with a 1, and a testing dataset with all number 2 values.

3. The Method 2 'K-fold Cross Validation Using Sites' training and testing datasets are created by first marking each hotel's usage records in the dataset with a unique number (e.g. there are 5 hotels, so the first hotel is numbered 1, the second hotel number 2 and so on). The first testing dataset is created using only the records of one hotel, with all other records in a training dataset. This is repeated until there are multiple testing and training datasets covering all combinations of hotels (e.g. the first testing dataset will contain only hotel 1 records, with the training dataset containing hotel records 2, 3, 4 and 5. The second testing dataset will contain hotel 2, with training dataset containing 1, 3, 4 and 5, and so on).

4. The Method 'K-fold Cross Validation Using Random Selection' training and testing datasets are created by marking each daily usage record row in the dataset with a randomly selected number between 1 to 4 (based on a 4-fold split). The first testing dataset is created based on the first random number's records, with the remaining records being used to create the training dataset. This is then repeated until there are multiple testing and training datasets covering all combinations of random values (e.g. the first testing dataset will contain only random number 1 records with the training dataset containing random records 2, 3 and 4. The second testing dataset will contain random number 2, with training dataset containing 1, 3 and 4, and so on).

5.4 Stage 4 - Profile Creation

The fourth stage uses the training datasets created in stage 3 to generate profiles that will be evaluated in stage 5. The aim of the profiles is to provide a representative daily usage pattern of the selected asset for a defined timeframe. For example, 12 profiles might be created that represent the average usage of hotel chillers for each month of the year. Each created profile will cover the expected usage levels of the chiller for each half-hour period over a day for the selected month, based on the sub-metered data from existing chillers. This stage covers profile creation by first describing the key background elements of the profile creation process that includes: (1) clustering method selection, (2) profile predictor options, (3) how the profile distributions are created by applying the clustered data and predictor selection, and (4) how values for the profiles are selected from the distributions. The outcomes from these elements are then used for a detailed breakdown of the technical steps taken to create the profiles.

5.4.1 Review of Clustering Methods

To create the profiles relies on first grouping by similarity the training data defined in stage 3. Grouping of time series data is normally undertaken using clustering methods as highlighted by Aghabozorgi et al. (2015) who identified over 20 options. In order to enable selection of an appropriate clustering method for this research, a review of existing literature was undertaken to identify which methods have previously been used for profile creation and determine their suitability for this research. Räsänen et al. (2010) looked at using self-organizing maps (SOM) in conjunction with K-means and hierarchical clustering methods for creating profiles from large amounts of customer-specific hourly measured electricity use data. These researchers found that combining SOM with k-means offered the best overall performance. Zhou et al. (2013) then looked at the usage of K-means, Fuzzy c-means (FCM), Hierarchical, and SOM clustering methods for electric load classification in a smart grid environment. These researchers found that while the traditional K-means, FCM and hierarchical methods work, more research is needed on innovative approaches that improve efficiency given the increasing sizes of the datasets being classified.

Chicco (2012) undertook an overview and performance assessment of the K-means, fuzzy K-means, follow-the-leader and hierarchical clustering methods for electrical load pattern grouping. He found that if the purpose of clustering is to extract outliers, then the hierarchical method is recommended. Whereas if the purpose is for grouping similar data, then the K-means method performs best. Panapakidis et al. (2014) informed the premise of the baseline comparison estimation method described in the previous chapter, and also reviewed the FCM, minimum variance criterion, SOM and K-means clustering methods in relation to pattern recognition algorithms for electricity load curve analysis of buildings. Its authors identified the SOM and K-means methods provided the lowest clustering error.

Based on the reviewed literature, it was determined that for the purposes of this research, the K-means clustering method would be used because: (1) the literature highlights that K-means clustering is a traditional, proven method used by researchers for time-series clustering, (2) literature shows that K-means has continued to perform well considering current clustering advances, (3) the literature also characterises the K-means clustering method as a simple, understandable, efficient and scalable clustering method.

Sayad (2017) describes how the K-means clustering method works by first selecting how many groups the dataset will be clustered into. For each group, a random point within the dataset is selected and deemed the centroid value. Each value in the dataset is assigned to the closest centroid. The mean of the values for each centroid is then calculated. The centroids are then moved to the mean position, and the values are reassigned to the now closest centroids. This process is repeated until a pre-defined number of interactions is achieved, or the level of centroid position change reaches a set tolerance. However, there are limitations noted in the literature that need to be considered when using the K-means method. The three main limitations and mitigations applied in managing them are:

- **Limitation 1** - It can be difficult to determine the Number-of-Clusters to use. In mitigation, a range of cluster values will be included in stage 6's input optimisation to determine the optimal cluster number to use for this dataset.

- **Limitation 2** - The initial cluster centre can significantly affect the outcomes. In mitigation, the K-means algorithm will be run multiple times with different centroid seeds to identify the final results with the best output of consecutive runs in terms of inertia (Scikit-learn, 2016).

- **Limitation 3** - It can be sensitive to noise and outliers. In mitigation, during review of the results of the stage 6 optimisation results, consideration will be given to identifying any abnormal results that might have been influenced by noise and outliers, and require additional analysis and correction.

5.4.2 Selecting Predictors for Profile Usage

The profiles created in this stage from the clustered data will be evaluated in stage 5 to estimate the chiller usage levels for the days contained within the testing datasets. Successfully applying these profiles requires a common identifying variable which links a profile to each day in the testing dataset. This identifying variable is known as the 'predictor value'. When using the profile method to estimate the DSR potential of a new site, only limited information will be available, namely the new site's business category (i.e. hotel) and the kW size of the DSR assets (i.e. a 200kW HVAC chiller). This limited information means that only date-based predictor values are available as a common variable between the profiles and a new site. For example, profiles could be generated based on the predictor being the week of the year. This then allows a DSR estimation to be completed by applying the profiles to the sites' DSR asset kW ratings over a year based on each week of the year. Date-based predictors have been used successfully by Van Wijk et al. (1999) for cluster identification of time-series data, and can then be utilised for further analysis. The example proposed in their research was for visualisation, yet can equally be applied as a predictor for DSR. Using only date-based predictors, the following four options were identified as variables for predictor values which will be included in stage 6's input value optimisation to determine the optimal date-based predictor type to use for this research's dataset:

- Week-of-Year: Weeks 1 to 52.
- Month-of-Year: Months 1 to 12 (January to December).
- Day-of-Year: Days 1 to 7 (Monday to Sunday).
- Weekend-Weekday: Saturday & Sunday or Monday to Friday.

5.4.3 Determining Weighted Distributions of Usage Values for each Profile

Once the data has been clustered and a predictor value type selected, this information then needs to be used to make a distribution of daily usage values, which can then be used to create the usage profiles. Each predictor value will be associated with multiple clusters. For example, if the Week-of-Year predictor is used, then the week 1 usage values may have been assigned to three different clusters with 60% of week 1 usage values being in cluster 1, 30% in cluster 2, and 10% in cluster 3. To create the week 1 profile first requires creating a weighted distribution of usage values from each associated cluster to ensure the profile has a representative distribution of likely outcomes. To create the distribution, a stochastic process is undertaken to randomly sample values from each cluster based on the percentage weightings. The unknown factor of this sampling is how many random values are required to create a representative distribution. An appropriate sample size is hard to determine as there is often a trade-off between good accuracy but high computational costs or poor accuracy with low computational costs (Royset, 2013). To determine the appropriate sample size involves testing varying sizes to understand the impact of variables on the final profile outcomes. Therefore, a range of sizes will be tested during stage 6's input value optimisation process. The outcomes of the optimisation process will be used to determine the stochastic selection sample size, with choices made based on assessing the computational requirements against accuracy outcomes.

5.4.4 Managing Non-Normal Distributions for Creating Profiles

The profiles will be created using the probability distributions generated using the stochastic process. To understand which statistical selection methods could be used to create the profiles requires first determining whether the probability distributions are normal. To assess if the distributions were normal, a test was undertaken using all the chiller asset sub-metered usage data to generate distributions based on using 5 clusters with the Month-of-Year predictor and 1000 random sampling points. This resulted in 576 distributions (12 monthly profiles multiplied by 48 half-hour periods per profile) each containing 1000 usage values. The distributions were then analysed used the Anderson-Darling and Skewness tests.

The Anderson-Darling test was used to determine if a data sample came from a population with a specific distribution. It is a modified version of the Kolmogorov-Smirnov test, in that it provides more weighting to the tails, and is recommended for large distributions, as used in this research (NIST/SEMATECH, 2012a). Using Python, the Anderson-Darling test was undertaken on each of the 576 distributions. The lowest result was 1.89, which was higher than all critical values including the 95% significance level threshold of 0.784. This means that the Anderson-Darling test shows that none of the distributions were deemed normal.

The Skewness test was used to assess that level of skew in the distribution (NIST/SEMATECH, 2012b). Using Python, the Skewness test was undertaken on each of the 576 distributions. The results of the test showed that at a 95% significance level, only 25% of the distributions p-value values were less than or equal to 0.05. This means that 75% of the distributions are skewed, with 86% being positively skewed and 14% negatively skewed.

Due to the distributions being non-normal and primarily positively skewed, the profiles' values cannot be obtained using mean and standard deviation values as these will not accurately represent the distributions. Instead, the middle profile value will be obtained using the median value, as this provides a more realistic middle point than the mean value (Laerd, 2017). The variance and spread within the profile will be measured using quartiles and 95% confidence intervals values. The upper and lower quartiles represent usage values that are +/- 25% from the median, and will be used to understand the level of variance around the centre. The 95% confidence intervals values are obtained by selecting as distributions the 2.5% and 97.5% percentile values, and represent the minimum and maximum ranges of potential usage with outliers being removed through rejection of the 2.5% of highest and lowest values in the distributions. By using these value ranges, it will be possible to understand the level of variances in the estimations, supporting more informed and reliable DSR assessments of new sites.

5.4.5 Creating the Profiles

The previous sections 5.4.1 to 5.4.4 defined the approaches that will be used to create the profiles using the training datasets defined in stage 3. To implement these approaches requires undertaking several process steps, and selecting appropriate input values. The key input values that need to be chosen to create the profiles are: (1) The number of clusters to create as outlined in section 5.4.1, (2) which profile predictor to use as outlined in section 5.4.2, and (3) selecting the number of random values to use for creating the profile creation distribution as outlined in section 5.4.3. As each of these inputs have multiple options an optimisation process is required to determine the optimal input values based on the dataset being used. Optimisation it is addressed separately in stage 6, due to its complexities. The following steps outline the profile creation process based on knowing which input variables are being used. The optimisation process in stage 6 will then use these steps to test all required combinations of inputs.

1. The first profile creation step requires clustering of the daily usage records in the training dataset using the K-means method based on the chosen number of required clusters. The K-means clustering method as described in section 5.4.1 was undertaken using Python with the results being used to update the training dataset with a new column that specifies which cluster each day has been assigned to.

2. The selected predictor input value is then used to understand how many daily usage records belong to each predictor variable, applying the following approach:
    a. Based on the profile predictor selected, an array is created that contains a column for each cluster and a row for each predictor value. (E.g. if the profiles are going to be based on Month-of-Year then the array will have 12 rows).
    b. For each predictor value row in the array, the clustered training dataset is scanned to identify all daily usage records for that predictor value. (E.g. for row 1 it will identify all January records).
    c. The number of records per cluster for each predictor is then counted. (E.g. the number of January records in cluster 1, cluster 2 and so on).
    d. Using these counts, the percentage split of cluster usage for each predictor is calculated and saved into the array. (E.g. a January predictor value might have a cluster split of 60% of daily records for cluster 1, 30% of daily records for cluster 2, and 10% of daily records for cluster 3).

3. For each profile predictor an array of half-hourly usage records is created, using the previous percentage-of-clusters array to determine a weighted random selection of values from each applicable cluster, and applying the following steps:
    a. An empty array of 48 columns (one per half-hour) with the number of rows equalling the number of random values to use for the stochastic selection process is created. (E.g. 1000 random values equals this many rows in the array).
    b. The training dataset is split into separate arrays based on the assigned cluster.
    c. Based on the percentage of cluster usage per profile predictor values, the number of random selections required from each cluster is calculated, and this many samples populated into the array from the required clusters. (E.g. if cluster 1 appeared 60% of the time in January, then this translates to requiring 600 random samples from this cluster).

4. The profile predictors half-hourly usage record arrays are then converted into the final usage profiles by taking the following steps:
    a. The following percentile values for each half-hourly column in each array are found:
        i. 2.5% - This represents the lower part of the 95% confidence level values of the usage records.

ii. 25% - The lower quartile boundary.
iii. 50% - The median value.
iv. 75% - The upper quartile boundary.
v. 97.5% - This represents the upper part of the 95% confidence level values of the usage records.

b. The percentile values for each profile are then saved as the output of this process.

5.5 Stage 5 - Profile Evaluation

Once the profiles have been created in stage 4, these then need to be evaluated against the testing dataset created in stage 3 to determine the level of estimation error. This section outlines the evaluation methods used before describing how the evaluation is undertaken.

5.5.1 Evaluation Methods

The evaluation of the profiles will use the same methods as were applied in the previous chapter 4, namely the Mean Absolute Percentage Error (MAPE) and Mean Bias Error (MBE). These methods remain valid as the profile evaluation is performed on the same data types, with the testing dataset representing the actual values that will be compared against the profile's estimated values. Using the same evaluation methods also allows for a direct comparison of the DSR estimation method outcomes from chapter 4 with the new profile method created by this research, as described in this chapter.

The MAPE method is applied by first calculating the difference between the forecasted and actual electricity usage, and then dividing this value by the actual usage for each half-hour period of the year. The absolute value of the outcomes is then summed and divided by the number of data points used. The resulting value provides a percentage indication of the level of fit between the forecasted and actual values, with lower being better. The MBE method is applied by first calculating the difference between the forecasted and actual electricity usage for each half-hour period of the year. The differences are summed and then divided by the sum of the measure electricity usage over the year. The resulting percentage shows the overall level of difference and bias between the forecasted and actual values. Further information on these methods, including benefits and limitations, is provided in section 4.1.3.1 for MAPE, and section 4.1.3.2 for MBE.

5.5.2 Undertaking the Evaluation

The evaluation of the profiles created in stage 4 is completed by testing one dataset at a time. If the data selection method produces more than one combination of training and testing data, then each set is evaluated separately (including a recreation of profiles in stage 4), and the results are then combined for analysis. The evaluation of a testing dataset and profile is undertaken by applying the following steps:

1. A dataset of estimated usage values is created used the new profiles. This is done by creating an estimated values dataset with the same size as the testing dataset. The estimated dataset is then populated by first matching the appropriate profile to each row in the dataset based on the predictor used. The half-hourly median values for each profile are then used to update the applicable rows in the estimation dataset. (e.g. if the row has a date of 2/1/2016 and the predictor is based on month, then this row is associated with the January usage profile).

2. The MAPE value is then calculated using equation (7) by first obtaining for each actual usage value in the testing dataset the corresponding estimated dataset value. The estimated value is subtracted from the actual value before being divided by the actual value. The resulting outcomes are turned into absolute values, before being averaged to determine the MAPE.

$$MAPE_{year} = \frac{\sum_{1}^{n}\left|\frac{A-F}{A}\right|}{n} \tag{7}$$

   Where:
   F = Forecasted value
   A = Actual value
   n = Number of values

3. The MBE value is then calculated using equation (8) by first obtaining for each actual usage value in the testing dataset the corresponding estimated dataset value. The actual value is then subtracted from the estimated value, with the resulting outcomes of all calculations being summed. This total is then divided by a sum of all actual values to determine the MBE value.

$$MBE_{year} = \frac{\sum_{1}^{n}(F-A)}{\sum_{1}^{n}(A)} \tag{8}$$

   Where:
   F = Forecasted value
   A = Actual value
   n = Number of values

4. If the data selection methods created more than one set of training and testing datasets, then the results of all the MAPE and MBE evaluations are averaged to obtain the final outcome.

5.6 Stage 6 - Optimal Method and Input Value Selection

Throughout the previous five profile creation stages questions have been raised regarding which development methods and input variables to use. The questions that need to be addressed can be split into two categories: (1) method development questions that once answered become part of the process for future profile creations, (2) recurring questions that need to be answered each time a new profile is created. A summary of the five questions raised during the previous stages and their categories are:

- Method development questions about the profile creation process:
  1. From section 5.2.1 – Which missing data option should be used?
  2. From section 5.3.1 – Which training and testing data selection method should be used?
  3. From section 5.4.3 – What stochastic sample size should be used?

- Recurring questions, for any new profile creation:
  4. From section 5.4.1 – How many clusters to use?
  5. From section 5.4.2 – What predictor value to use?

The optimisation stage process is covered in four sections. Section 5.6.1 covers selection of the input values and generation of the optimisation results. Section 5.6.2 uses the optimisation results to answer the three method development questions about the profile creation process. Section 5.6.3 uses the optimisation results to address the two recurring questions for any new profile creations. Section 5.6.4 summarises the outcomes for each question in the context of the data analysed for this research.

5.6.1 Generate Optimisation Results

Parametric programming will be used to generate the optimisation results that will be analysed in the sections 5.6.2 and 5.6.3. Parametric programming is described by Pistikopoulos, Galindo, & Dua (2007) as aiming to '*obtain the optimal solution as an explicit function of the parameters*', and is based on sensitivity analysis theory. The authors describe the main difference between the two theories by stating that parametric programming provides a complete map of the optimal solution, while sensitivity analysis provides solutions in the neighbour of the nominal values. Parametric programming offers different implementation methods, depending on if there are computation resource constraints that require pre-processing to reduce input value ranges, or if sufficient computations resources are available to enable all input variations to be tested. Based on running initial tests of the required input parameters for this research, it was determined that there was sufficient computation power available to undertake a full assessment of all input values without pre-reduction. This enables a complete mapping of potential outcomes for analysis.

The first step to generating the optimisation results is to define what input values will be used. The five questions that are to be answered through this process each have a single input with multiple potential values. Each question is analysed below to determine what values are going to be used, and the selected values being listed in Table 30. Once the optimisation process has been completed the results will be analysed in sections 5.6.2 and 5.6.3 to determine the answers for each of the following questions:

1. **Which missing data option should be used?** This question was raised in section 5.2.1 and aims to understand the impact of fixing missing data through interpolation. This is a method development question that once answered in this section will determine the recommended approach that will be used for managing missed data in future profile creations. Table 27 shows that there are 63 days out of the 4012 days of data that have at least one missing half-hourly usage reading. As the profile method uses whole days of usage information, any day containing missing values will be excluded unless fixed. Based on the various levels of missing values in Table 27 four missing data options were selected for assessment to understand their impact on error levels: (1) exclude all 63 days containing missing values, (2) fix the 29 days that contain 1 missing value, (3) fix the 51 days that have between 1-6 missing values, and (4) fix all 63 days by addressing missing values of between 17-25.

2. **Which training and testing data selection method should be used?** This question was raised in section 5.3.1 and aims to understand which of the three data selection training and testing data selection options should be used. This is a method development question as the selected method will be recommended for use in future profile creations. There will be three values for this input that cover each method as described in section 5.3.1: Out of Sample, K-fold Cross Validation Using Sites, and K-fold Cross Validation Using Random Selection.

3. **What stochastic sample size should be used?** This question was raised in section 5.4.3 and aims to understand how many random points are needed to be selected from each cluster when generating the dataset used to create each profile. This is a method development question as the result of the optimisation will determine the recommended sample size for use in future profile creations. Two value ranges were selected for this input, the first covers a lower range of values from 100-1000, while the second ascertains the impact of larger values from 2000-10000.

4. **How many clusters to use?** As outlined in section 5.4.1 the K-means clustering method requires as an input the number of clusters to be used. As future profile creations will use different datasets, this optimisation question needs to be answered each time the process is run. Therefore, to determine the optimal number of clusters to use for creating the profiles this input is including in the optimisation process. While the number of clusters could range from 1 to 1000s, the final range selected for the optimisation process was from 1 to 10 based on initial tests showing that the results plateaued after the 3rd cluster, as seen in Figure 20.

5. **What predictor to use?** Section 5.4.2 identified four potential predictors that could be used to apply the profiles. As the optimal predictor needs to be selected for each new profile created, this question needs to be addressed each time. The values that will be tested each time for this input are: Week-of-Year, Month-of-Year, Day-of-Year, Weekend-Weekday.

Table 30 - Optimisation Values Used Per Input

Values for Method Development Questions			Values for Recurring Questions	
Missing Data Methods	**Data Selection Methods**	**Stochastic Sample Sizes**	**Profile Predictors**	**Number of Clusters**
Exclude All	Out of Sample	100	Week-of-Year	1
Fix gaps of 1	K-fold Cross Validation Using Sites	200	Month-of-Year	2
Fix gaps up to 6	K-fold Cross Validation Using Random Selection	300	Day-of-Year	3
Fix gaps up to 25		400	Weekend-Weekday	4
		500		5
		600		6
		700		7
		800		8
		900		9
		1000		10
		2000		
		3000		
		4000		
		5000		
		6000		
		7000		
		8000		
		9000		
		10000		

Having defined the optimisation input values in Table 30 the next step requires processing all combinations of values to generate the results dataset. This processing is undertaken using Python to create and evaluate profiles using the steps outlined in stages 1-5. A looping process creates and evaluates profiles using all combinations of values, with the MAPE and MBE values of the profiles being recorded. Based on the values in Table 30, this results in 9,120 combinations (4 Missing Data Methods x 3 Data Selection Methods x 19 Stochastic Sample Sizes x 4 Profile-Predictors x 10 Number of Clusters). For example, the initial values for the first iteration would be Missing Data Methods = *Exclude All*, Data Selection Methods = *Out of Sample*, Stochastic Sample Size = *100*, Profile Predictor = *Week of Year*, Number of Clusters = *1*. When this first iteration is complete and MAPE and MBE values obtained, then the number of clusters value is updated and repeated until all 10 values have been evaluated for this combination of values. Once all values in this loop are completed, then the next layer above is updated, i.e. the Profile-Predictors value is updated to *Month-of-Year* value and then reiterates through all 10 number of cluster values. This sequence of changes occurs until all 9,120 combinations of input values have been evaluated. The processing of 9,120 combinations took approximately 22 hours to run on a PC with a Solid State Hard drive, i7 CPU and 16GB RAM.

5.6.2 Use Optimisation Results to Answer Method Development Questions

In this section, the optimisation results are used to determine the optimal outcomes for each of the method development questions, namely: (1) which missing data method should be used? (2) which training and testing data selection method should be used? (3) What stochastic sample size should be used? As each of these questions relates to the method development they will be answered in the order they were identified to reflect that the answer will define the outcomes of that step which then feeds into the next question. To reflect the impact of the answer on the next question the optimisation results are reduced accordingly to only contain the outcomes that would have occurred if that answer had been used at the start. This ensures each question only uses relevant data.

5.6.2.1 Determining Optimal Missing Data Option:

To understand the optimal option for managing missing data, the interpolation method as described in section 5.2.1 was used to fix four different missing gap sizes. To understand the impact of each fix, the average MAPE values of the optimisation results were calculated as shown in Table 31. The averages show that fixing single missing gaps provides the best MAPE reduction of 0.29 percentage points compared to not performing any gap fixes. By fixing the single missing gaps, 29 days of additional data can be used, increasing the overall dataset size by 0.83% (due to any none fixed missing gaps causing the whole day's data to be excluded). Table 31 also shows that while increasing the fix beyond a single gap has a lower MAPE than not fixing any gaps, the actual improvement decreases as the length of gap fixing increases. This reflects the test performed in section 5.2.1, which shows the error level increases as the gap fix length increases. Based on these results, the optimal method for managing missing data using interpolation is to only fix single gaps. Using this option reduces the optimisation results from 9,120 to 2,280.

Table 31 - Optimisation Results Average MAPE for each Missing Data Method Selection
Based on 63 days out of the 4012 days of data containing at least one missing half-hourly usage reading that if not fixed is removed from the dataset

	Average MAPE	Days Fixed	Days Removed	Percentage of Dataset Removed
Exclude all missing values	59.25%	0	63	1.57%
Fix missing values with 1 half-hourly gap	58.96%	29	34	0.85%
Fix missing values with 1 to 6 half-hourly gaps	58.99%	51	12	0.30%
Fix missing values with 1 to 25 half-hourly gaps	59.06%	63	0	0.00%

5.6.2.2 Determining Optimal Testing and Training Data Selection Method:

The second method development question was raised in section 5.3.1 about which training and testing data selection method should be used. To answer this question the MAPE and MBE values from the updated optimisation results were used to create Figure 18 (see Table 50 in Appendix E for figure numbers). This box plot shows the minimum, lower quartile, median, upper quartile, and maximum values for each of the three methods. Figure 18 and the previously discussed background in section 5.3.1 informs the following discussion on the suitability of each method, and the final selection of which one should be used for evaluation of the profiles:

- **Method 1 - Out of Sample** – This selection method has the lowest average MAPE error level, but the largest spread of primarily negative MBE values by a wide margin in comparison to the other methods as shown in Figure 18. The negative MBE values indicate that the majority of estimated results are less than the actual values. The MBE outcome highlights the limitation of this method, as only using a testing dataset of the latest time segment can result in the evaluation being influenced from events that have occurred in that selected time segment, which might not provide a true reflection of the overall dataset. This suggests that this method should be rejected.

- **Method 2 - K-fold Cross Validation Using Sites** – This method has the highest average MAPE error level with a low spread of MBE values and an overall negative bias as shown in Figure 18. The previously noted limitation of this method is that testing each site individually risks the possibility that one of the sites might have a usage pattern different enough from the other sites to cause the error results to be skewed. For this analysis, the individual results of each site's tests were examined, which identified that the Worsley Park site was generating MAPE values that were consistently 20%-25% higher than the other tested sites. This demonstrates how one site can influence the overall results, which suggests that this method should also be rejected.

- **Method 3 - K-fold Cross Validation Random Selection** – This method had an average MAPE error level in-between the other two methods, and a low spread of MBE values that were mainly negatively bias as shown in Figure 18. This method's usage of random sampling means that it is less susceptible to the selection limitations noted in the other two methods as shown with its in-between average MAPE value. These factors suggest that this method could be acceptable for data selection.

Based on the analysis of each method's outcomes it was found that the K-fold Cross Validation Random Selection Using Sites and K-fold Cross Validation Using Random Selection methods achieved similar results. As only one method could be used for the final profile creation stage, the K-fold Cross Validation Using Random Selection was selected because: (1) it provides a balanced outcome with a MAPE median of 61%, which is in-between the other two methods medians of 59% and 63%, (2) the MBE values are primarily negative, which for DSR is preferred to avoid overestimation (as reviewed in section 4.1.3), and have a lower spread than the Out-Of-Sample method, (3) it is less susceptible to being influenced by any single block of usage data, a factor which could adversely affect the K-fold Cross Validation Random Selection Using Sites method if one site had any major differences in usage data, and (4) it is a recommended approach for time series forecasting verification (Bergmeir et al., 2015).

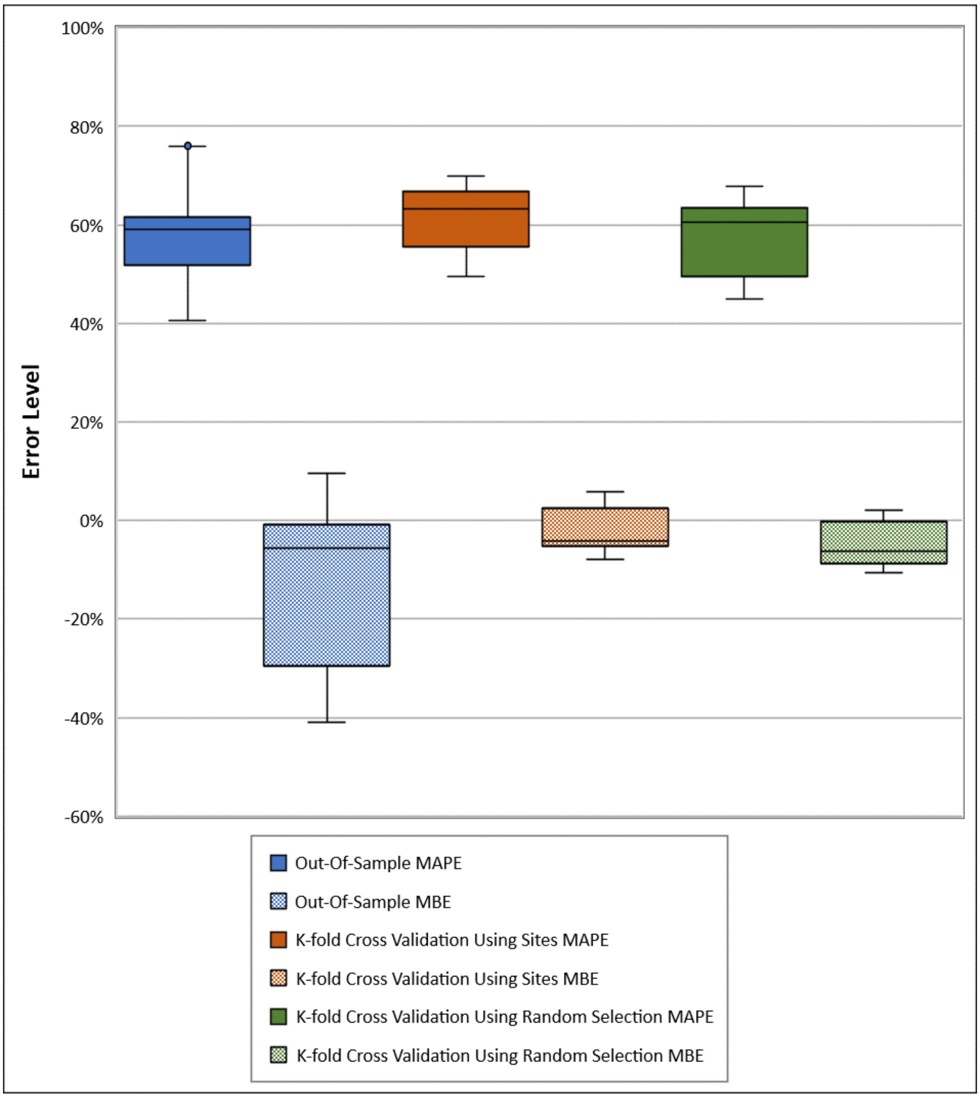

Figure 18 - Box and Whisker Plot Results of each Data Selection Method

5.6.2.3 Determining Optimal Stochastic Sample Size

The stochastic sample size input is the last method development question and was identified in section 5.4.3. As noted in section 5.6.1, Royset (2013) points out that selecting the stochastic sample size is a balance between computational costs and accuracy. This is demonstrated by the optimisation process for this research taking approximately 22 hours to run, with the time taken to process each iteration increasing as the stochastic sample size increased. To understand the link between processing time and the average MAPE values, the time taken to complete processing of each stochastic sample size value was calculated and then added to Figure 19's secondary axis.

Figure 19 shows that the average MAPE value only improved by 0.17% when increasing the stochastic sample size from 100 to 10,000 (see Table 51 in Appendix E for figure numbers). The largest improvement occurred during the initial increase from 100 to 400, before a smaller steady improvement until the sample size hits 3,000, at which point the results plateau. While increasing the size only provided a marginal MAPE improvement it did have a major impact on the time taken to process. Using 10000 points resulted in computational processing that took 42 minutes to complete all iterations using that value, compared with only 6 minutes to process when using 100 points. While the sample size impacts processing times, this impact is mitigated when the selected value becomes the default for future profile creations. As the difference in MAPE values is very small, a small sample size with a shorter processing time will have limited impacts on accuracy. Therefore, for the purposes of this research, the default sample size will be 3,000 as the next point (4000) decreases the MAPE by 0.0029% at a cost of increasing processing time by 22.7%. Using this value reduces the optimisation results from 760 to 40.

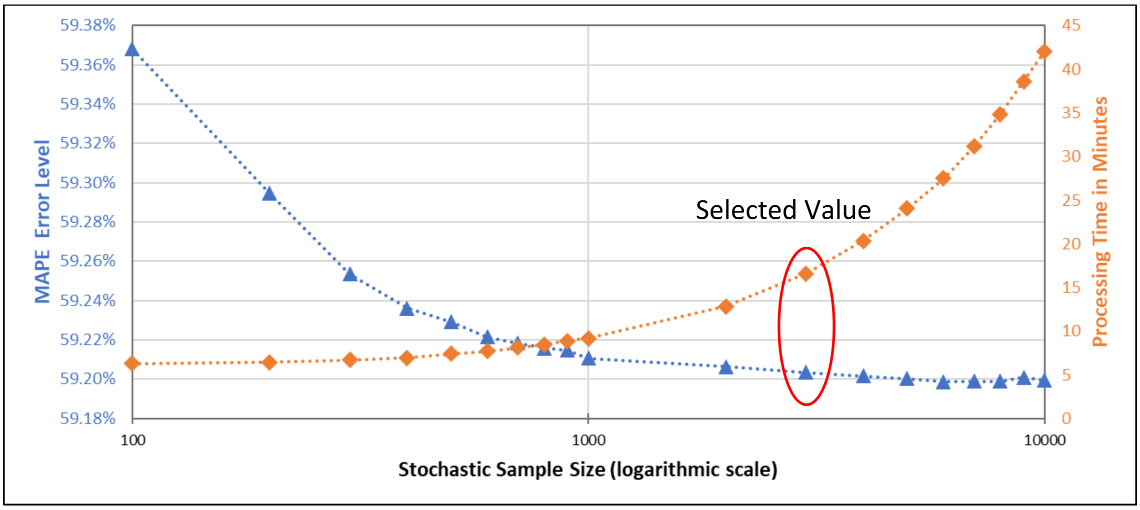

Figure 19 - Comparison of MAPE Versus Processing Time for The Stochastic Sample Size Selection

5.6.3 Use Optimisation Results to Answer Recurring Questions

Addressing the method development questions in section 5.6.2 has reduced the optimisation results from 9,120 to 40. The previous method development questions were answered in the order that they occurred in as the outcomes impact the next process step (for example, missing data is managed in stage 1 and then used during stage 3 data selection). In contrast, the recurring questions do not have a predetermined order as they are interlinked in impacting the final outcome. Based on the optimisation results containing all possible outcomes, two methods will be used to answer the recurring questions with the outcomes being compared to determine how they impact the final input selection. The two methods are described by May et al. (2011) as Exhaustive Search and Backward Elimination. The Exhaustive Search method works by evaluating all possible combinations and then selecting the optimal combination based on predetermined criteria. For this research, applying this method would involve finding the combination of inputs that produce the lowest MAPE value. This method has the benefits of providing a complete search of all options, a clear output of an optimal solution, and can be fully automated. Yet it is limited by not providing a comprehensive view of how the values were selected due to the automatic selection of the inputs that create the lowest MAPE value. The automatic selection also causes a potential risk of selecting a set of outlier values that while providing the lowest MAPE value might be sensitive to small input values changes causing the MAPE values to increase substantially.

The second method, Backward Elimination, also works by evaluating all possible combinations, but uses a different approach towards selecting the optimal combination. Once all combinations have been processed, the results are assessed to understand the level of impact each input has on the final outcomes. The impacts are used to manually order the input values and then the user assesses the lowest impacting input first, in order to decide which value should be used as the optimal value. The selection process then continues until all inputs have been assessed in ascending order of impact. This method has the benefits of providing a complete search of all possible options, a greater understanding of why input values were selected, and lower risk of selecting outliers. Yet its limitations include potentially being difficult to determine which order to assess the inputs if during the assessment of the optimisation outcomes there is not a clear range of impacts across the inputs (though this can mean that any order could be used), it can be difficult to automate, and relies on user interpretation which could lead to mistakes.

5.6.3.1 Backward Elimination Determining of Analysis Order

The first step in using the Backward Elimination method requires determining the impact of each input so that the order of analysis can be determined, starting with the input that has the lowest impact. Determining the impact and order was achieved by averaging the MAPE values for each input based on the remaining 40 optimisation results (e.g. each of the ten cluster values multiplied by 4 predictors values). The difference between the maximum and minimum MAPE averages for each input was then determined and used to assess the overall impact of each input to decide the order for analysis. This resulted in the following average MAPE ranges and order for the Backward Elimination process that will be undertaken over the next two sections:

1. Number-of-Clusters MAPE range: 6.26%.
2. Profile-Predictor MAPE range: 11.71%.

5.6.3.2 Backward Elimination Selection of Number of Clusters Value

The ‘number of clusters’ input as described in section 5.4.1 had the lowest impact, with a MAPE range of 6.26%. To determine the optimal number of clusters, the 40 optimisation results are split into 10 groups based on the cluster value used. For each group, the average MAPE results were calculated and plotted in Figure 19 (see Table 52 in Appendix E for figure numbers). The chart shows that the error level reduces until cluster value 3, after this the reduction plateaus with only minor reductions in errors until the lowest level is archived at cluster 5. The difference between cluster 3 and 5 is 0.03 points going from a MAPE of 49.42% to 49.39%. Based on this analysis, the optimal cluster size for this profile creation for this dataset is 3.

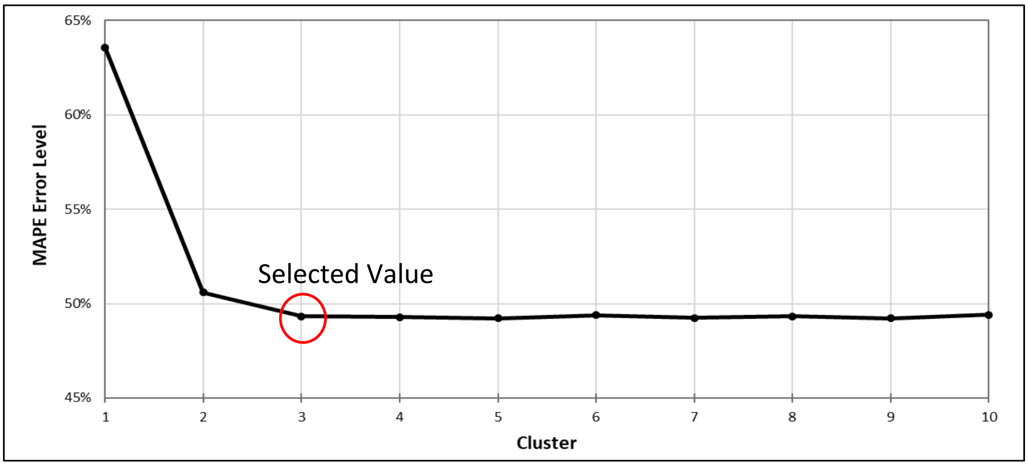

Figure 20 - Comparison of Number-of-Clusters Input Values

5.6.3.3 Backward Elimination Selection of Profile Predictor Value

The final input that had the highest impact was the profile predictor, with a MAPE range of 11.71%. The determination of this variable's optimal input would normally use a subset of the optimisation results based on a cluster size of 3. However, as with the previous variable, this research will not limit the data selection for the purposes of enabling an understanding about how the values would compare if this restriction was not in place. Therefore, the MAPE results for each cluster of each predictor value was plotted in Figure 21 (see Table 53 in Appendix E for figure numbers). The chart shows that the Month-of-Year predictor value performs the best with Week-of-Year being a very close second with a MAPE error only 0.51 points higher. The overall results show only minor variations based on cluster usage, which indicates the outcome of the profile prediction would not be impacted if another cluster size or site selection had been chosen.

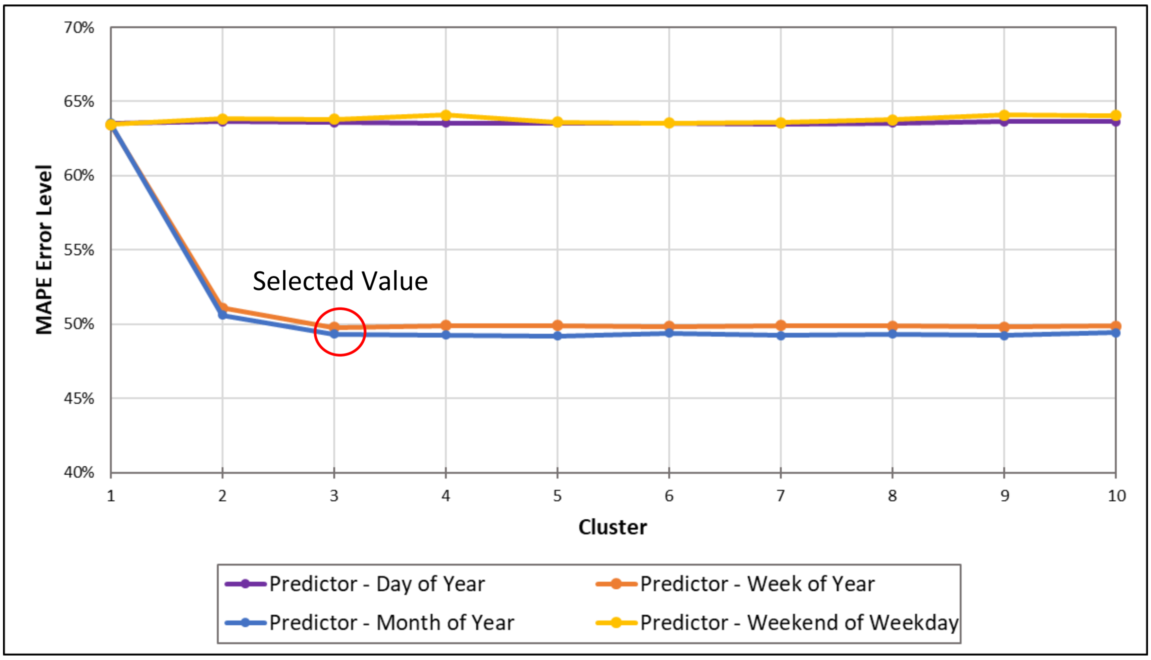

Figure 21 - Profile-Predictor Comparison

5.6.3.4 Exhaustive Search Selection of Values

The Exhaustive Search method is the other approach that can be used to finding the optimal input values. It works by finding the combination of input values that generates the lowest overall MAPE value. To find these values, the MAPE results from the 40 combinations of input values were ranked in ascending order. The combination of inputs with the lowest MAPE value of 49.39% was then identified based on a cluster size of 5 and used the month-of-year predictor.

5.6.3.5 Comparison of Backward Elimination and Exhaustive Search Methods

The outcomes of the Backward Elimination and Exhaustive Search methods are recorded in Table 32 to allow for comparison. The results show that there is only a 0.03 point MAPE difference between the two methods. This difference is the result of the Exhaustive Search automatically selecting the 5 cluster input value that produces the lowest MAPE value. Whereas the Backward Elimination as reviewed in section 5.6.3.2 selected the 3 cluster input value as this was the point at which additional clusters only produced margin differences in MAPE. Based on the similar outcomes of both approaches this comparison indicates that either method would be acceptable for selected input variables during each profile creation. For the purpose of the hotel chiller profiles being created in this research the Backward Elimination values will be used as the manual nature of this method provide more informed selection then the Exhaustive Search method.

Table 32 - Comparison of Exhaustive Search and Backward Elimination Input Values Selection

Value	Exhaustive Search	Backward Elimination
MAPE	49.39%	49.42%
Number of Clusters	5	3
Profile Predictor	Month-of-Year	Month-of-Year

5.6.4 Summary of Optimal Method and Input Value Selection

This stage 6 seeks to answer questions raised during the development stages 1-5 through usage of optimisation methods that enable all options to be assessed when deciding the optimal outcomes. This process encompassed an extensive amount of analysis over sections 5.6.1 to 5.6.3.5, therefore Table 33 provides a summary of the questions raised and answers provided, and will be used in the next stage 7 to finalise the profiles' development for the hotel chiller assets.

Table 33 - Summary of Optimisation Questions and Answers

	Questions		Answers	
	Section	**Question**	**Section**	**Answer**
Method	5.2.1	Which missing data option should be used?	5.6.2.1	Use the 'fix single gaps' option
Method	5.3.1	Which training and testing data selection method should be used?	5.6.2.2	Use the 'K-fold Cross Validation Random Selection' method
Method	5.4.3	What stochastic sample size to use?	5.6.2.3	Use a sample size of 3000
Recurring	5.4.1	How many clusters to use?	5.6.3.2	For this profile used 3 clusters
Recurring	5.4.2	What predictor to use?	5.6.3.3	For this profile use Month-of-Year predictor

5.7 Stage 7 - Final Profile Creation and Existing Method Comparison

This final stage of the development process utilises the results from stage 6 to create the final profiles that will be used to generate DSR estimations for comparison against the outcomes of chapter 4's analysis of four existing DSR estimation methods. To create the final hotel HVAC chiller asset profiles the values from Table 33 are used, and the following subset of the steps previously outlined in stages 1-4 are then applied:

- **Stage 1** – Sub-meter kW usage data for an asset type from a range of similar business is obtained. This research obtained HVAC chiller usage data from five hotels in the UK that will all be used for generating profiles.

- **Stage 2** – The sub-meter kW data is converted into half-hourly kW usage periods with any single missing half-hourly gaps being fixed using interpolation. The half-hourly values are then normalised by converting the kW usage value into a percentage of the assets maximum kW usage.

- **Stage 3** – For creation of the final profiles the data selection stage is not required as the full dataset of all sub-meter data will be used.

- **Stage 4** – The monthly hotel HVAC chiller profiles were created using 3 clusters and the Month-of-Year profile predictor with each profile being formed of the following percentile values: upper and lower 95% confidence levels, upper and lower quartiles, and the median.

The created profiles are now ready to be used for estimating the DSR potential of HVAC chiller systems in other UK hotels. To demonstrate the application of these profiles and enable comparison with the existing DSR methods, the newly created profiles will be applied to the same dataset used in chapter 4, as described in sections 5.6.4 and 5.7.2.

5.7.1 Generating DSR Estimations to Compare against Existing Methods

Chapter 4 compared four existing DSR estimation methods by applying each method to two years of data from two hotels. The resulting MAPE and MBE outcomes were compared to understand each method's level of error. Additionally, the cost of running each method was calculated and used to compare against the error levels to determine the links between cost and error. To compare the new profile method against the existing DSR methods the same approach is applied, treating the new profile method approach as DSR estimation method 5. To obtain DSR estimations from the profile method which are comparable with the outcomes described in chapter 4, new usage profiles are created using the optimal input values identified in stage 6 but excluding the sub-metered data of the hotel being estimated. These profiles are then used to create DSR estimations for the excluded hotel. This provides a realistic application of the profile method working on the basis that the hotel being estimated has not yet been enabled for DSR, and therefore no sub-metered data would be available. To generate the DSR estimations and resulting MAPE and MBE values the following steps were repeated twice to cover each hotel being analysed:

1. The steps in stages 1 and 2 were undertaken to create a normalised dataset from all five hotels' HVAC usage data, including the hotel that the DSR estimation is being undertaken for.

2. The stage 3 training and testing selection steps are modified so that the hotel being estimated is excluded from the training dataset, and instead its data is used to create two testing datasets covering the years used in chapter 4 (which were Regents Park, 2015 and 2016; Bristol Royal, 2013 and 2016). The remaining hotel data is used in the training dataset.

3. Stage 5 is undertaken using 3 clusters and the Month-of-Year predictor as determined in stage 6 to create the profiles. The profiles are then evaluated using stage 6 to calculate the MAPE and MBE values.

The resulting MAPE and MBE values for the profile method are then compared to the chapter 4 equivalent outcomes as reviewed in the results and discussion section 5.8. The final element of the comparison is to compare the profile method's cost of usage against the existing methods.

5.7.2 Profile Method Cost Calculation

The second element of chapter 4's comparison looked at how each method's usage costs compares with their MAPE error levels. To perform the same comparison for the profile method involves calculating the cost of applying it. To determine the costs involved, the information input costs table format from section 4.1.5 has been replicated in Table 34, and updated to show the two expected costs of the profile method only. The first cost is the same as in section 4.1.5, and covers the time required to contact the site to obtain the maximum kW rating value for the chiller assets and to run the profile method. The second cost is a share of the time over a year that might be required for a person to maintain and update the profiles as new data is obtained. This work has been estimated at requiring 1 day per week. However, it could be lower depending on the level of automation applied. Based on these two costs, it is estimated that each usage of the profile method will cost £26 based on 500 uses per year. This cost and MAPE values are compared against the outcomes of chapter 4 in the results and discussion section 5.8.

Table 34 - Information Input Costs for the Profile Method

Information Input	**User Time to obtain/use (minutes)**	**User Time cost (@ £20 per hour)**	**External Information Cost**	**Cost of External Information usage (@500 uses year)**	**Total Input and Usage cost**
Site Information gathering and applying profiles.	30	£10	Free	£0	£10
Updating or creating Profiles when new sub-metered data is available.	24000 per year	£8,000 per year divided by 500 uses = £16 per use	-	-	£16 per use
Total Method Cost					**£26**

5.8 Results from Developing the Profile Method

The results of the profile development are reviewed and discussed over the next four sections. Section 5.8.1 reviews the generated hotel chiller profiles as an example in order to formulate an understanding about what meaningful information profiles can provide regarding how usage varies during each profile day and between each profile. Section 5.8.2 looks at how the usage profiles help reduce uncertainty during DSR estimation. Section 5.8.3 then reviews sensitivity of the key inputs. Finally, section 5.8.4 compares the outcomes of the profile DSR estimation method with the four DSR estimation methods reviewed in chapter 4, to enable conclusions about whether it succeeds in reducing the uncertainty of DSR potential estimation when compared with other methods.

5.8.1 Review of Generated Profiles

Figure 22 and Figure 23 show the results of the twelve profiles generated in section 5.7 for hotel chillers (see Table 56 through Table 59 in Appendix E for figure numbers). These twelve profiles have emerged because of the method's use of the Month-of-Year predictor, as outlined in section 5.6.4. Each profile consists of six lines, with each line being formed from the 48 half-hourly percentage values of expected chiller load over a day. The upper and lower 95% confidence levels, upper and lower quartiles, and the median lines were calculated as part of the profile creation process outlined in section 5.4.5. The 'daily average' line was added to support analysis of the profiles in this section. Note that while the lines are connected in the profiles to facilitate trend analysis, the half-hourly values are actually independent of each other due to being selected from the usage distributions based on percentile values. This means there is no 'memory' between half-hourly values (i.e. the previous value does not directly influence the next value), and instead the profile lines only show the trend that the next value is likely to be in a similar value range. This is suitable for DSR estimation of a new site over a year as the load profiles can provide an indication of overall potential. However, the profiles are unsuitable for predicting the usage of a chiller on a specific day, as this would require different analysis to take into account additional forecasting variables (for example, outside temperature and current load levels) to determine short-term expected usage.

The median line will be used as the default values for examining the profiles because it represents the middle point of expected usage. Use of the median values also reflects the initial desktop DSR estimation used by DSR aggregators for a new site, as outlined in section 3.1.1 and 4.1. Given that desktop estimation occurs early in an onboarding process, the median line offers an initial indication of the site's DSR potential, which the aggregator can then utilise when deciding if the site should progress to the next stage – i.e. performing a site survey, where initial estimates will be refined. The implications of using the median line for DSR estimation will be further discussed in section 5.8.2. However, in this section, the profiles are examined to provide an overview of three major features of the generated hotel chiller profiles: (1) unusual load spikes, (2) variation across months, and (3) variations across days. These chiller profiles are then compared to existing research about end-usage load profiles. This information forms the basis for an in-depth look at section 5.8.2 into how the profile method's outputs could impact DSR estimation.

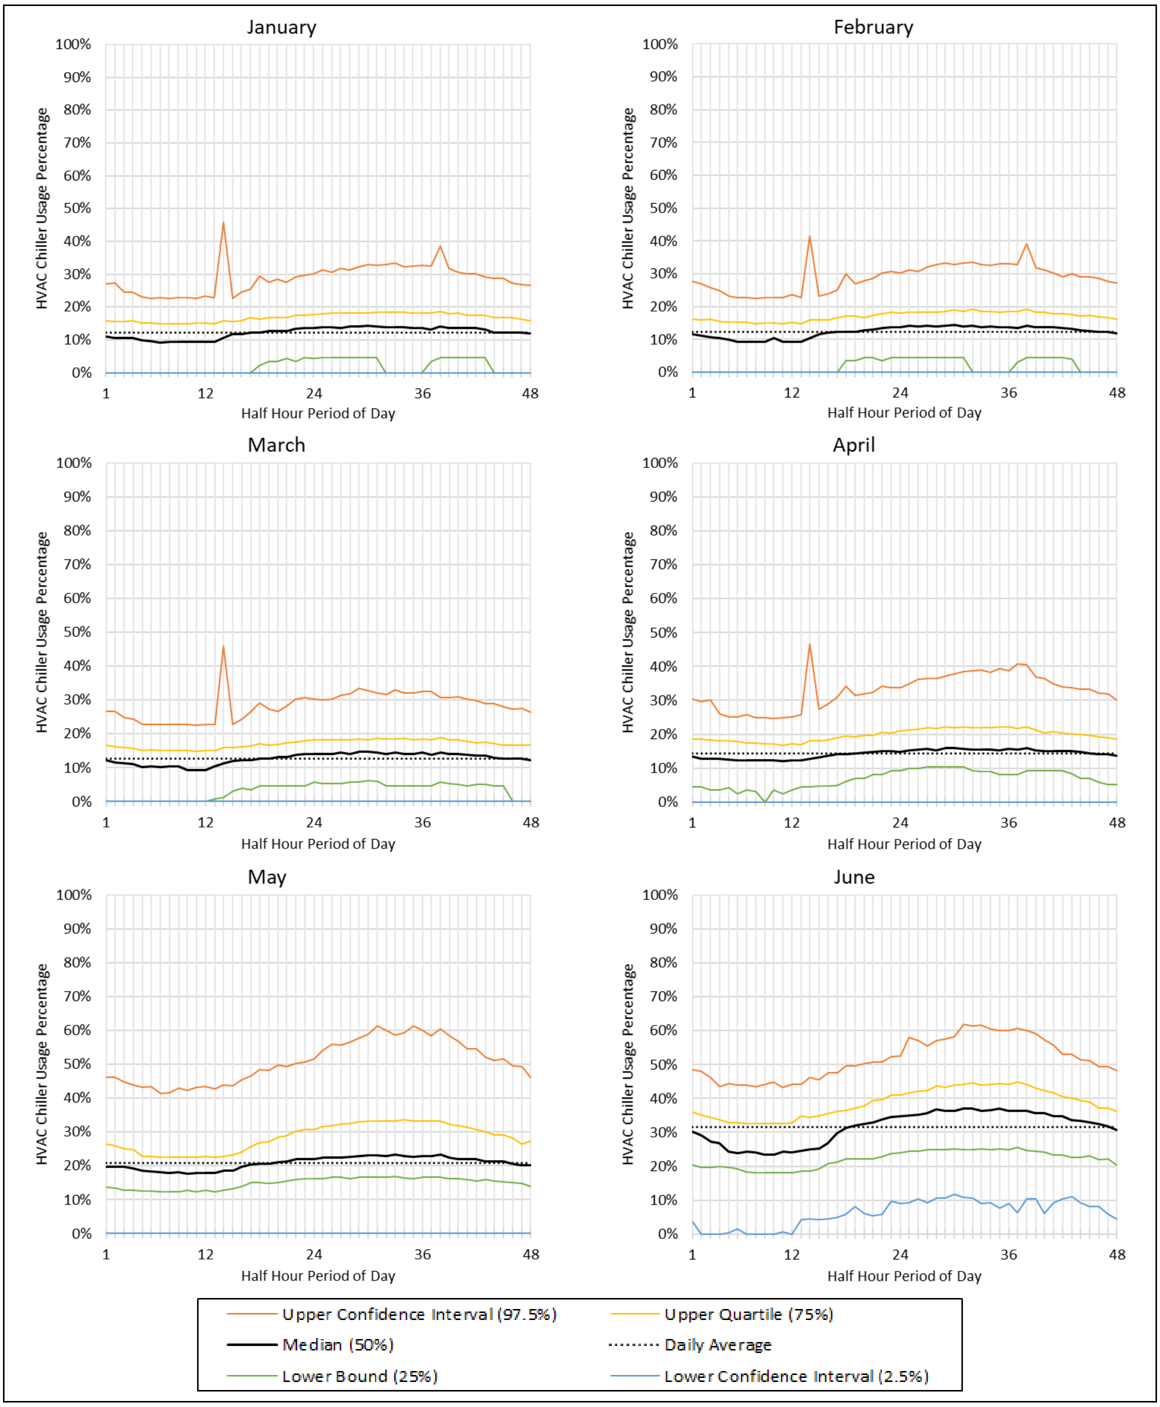

Figure 22 - Monthly Profiles Generated for Hotel Chiller - January to June

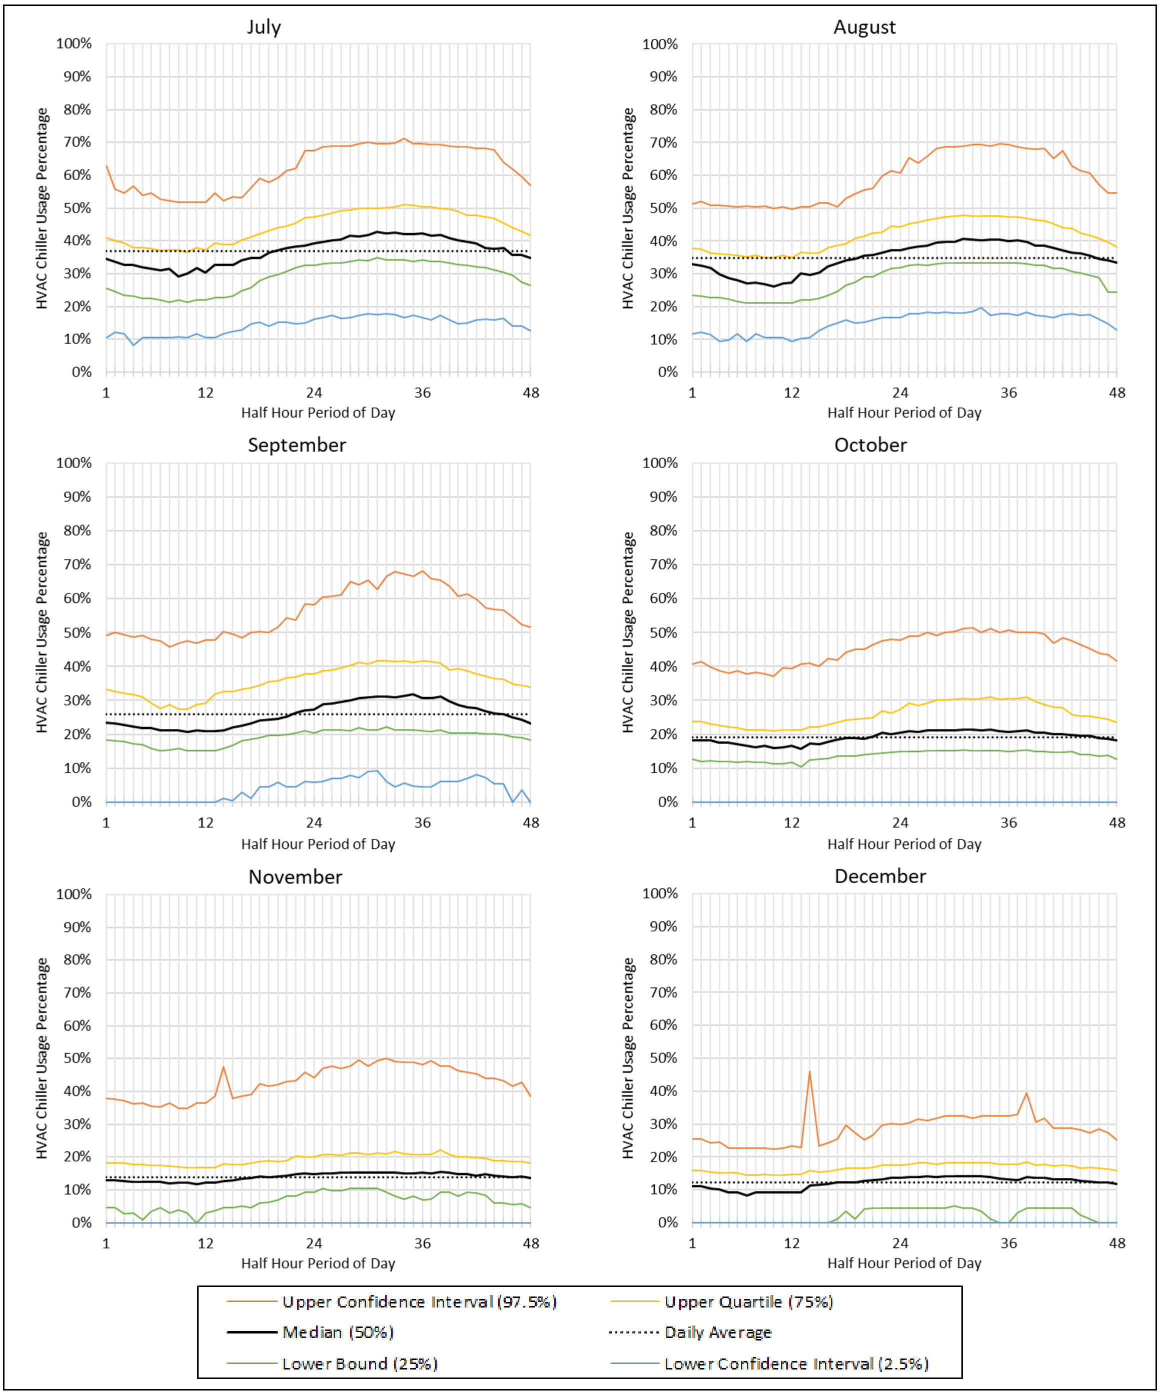

Figure 23 - Monthly Profiles Generated for Hotel Chiller - July to December

The first discernible major feature of the profiles in Figure 22 and Figure 23 is the apparent load spikes, which occur in the upper confidence interval lines at 06:30-07:00 (half-hour period 14) from November to April, and 18:00-18:30 (half-hour period 38) from December to February. The half-hourly usage data that forms these profiles attributes these spikes to the Maida Vale hotel chiller usage patterns. The data for the morning spikes shows that over this period the hotel turned off their chiller from midnight until 06:30 on 87 days. When the chiller was turned on again, the usage levels for the half-hour period of 06:30-07:00 increased to between 40.7% and 57.8%, before dropping to between 13.5% and 22.4% after 07:00. Figure 24 illustrates this behaviour and shows the difference between profiles with and without memory effect by plotting a new line onto the existing December load profile of actual usage of the HVAC chiller at the Maida Value hotel on the 23/12/2013 (a Monday). This actual usage line shows how the spikes are directly related to this hotel's HVAC usage and how after being turned on the usage levels drop to a steady state that matches the medium level usage. This initial start-up usage spike can be attributed as being caused by the chiller having to bring the system cooling fluid back to within a pre-set temperature range. The memory effect can be observed using Figure 24, as the actual usage line shows how the previous half-hour's usage level will influence the likelihood of the next half-hour, except when manual control is used to turn the chiller on or off (see Table 60 in Appendix E for figure numbers). In contrast, the load profiles created for December do not show memory effect, which can be seen from the actual usage line jumping between the lower conference, upper conference, and medium levels across the day.

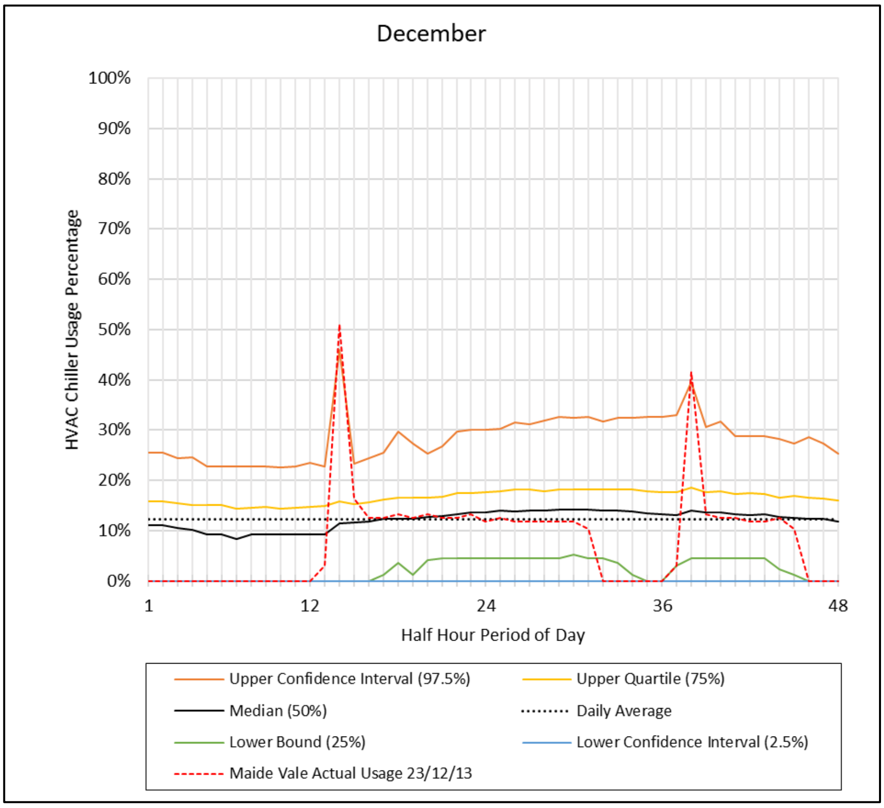

Figure 24 - Comparison of Actual Usage Against Load Profile

The afternoon spike between 18:00-18:30 occurred after the chiller had been turned off between 16:00 to 18:00 over 63 days, resulting in peaking of between 30.3% and 49.6% before dropping back to between 11.2% and 20.5% after 18:30 as illustrated in Figure 24 This spike reflects the 'rebound' effect, which can be experienced during DSR events whereby the turned off asset causes higher usage after the event as it returns the system back to its pre-event state (Palensky & Dietrich, 2011). This turning off pattern also matches businesses that participate in 'Triad Avoidance'. The National Grid recognises this practice as a method to reduce a business's Transmission-Network-Use-of-System-Charges (National Grid, 2016d). Businesses are charged for using the transmission lines based on the amount of electricity used during the three highest system demand half-hour settlement periods between November to February, in each case as determined by the National Grid. By having a lower usage during these times, businesses can reduce this cost. However, users can only guess when these three periods might occur based on past experience. Data shows that these normally occur between 16:00 to 18:00, and that all three occurred between 17:30 to 18:00 during the 2016/2017 November to February period (National Grid, 2017g).

The second noticeable major feature from the profile method's twelve HVAC chiller profiles is the variance in usage levels between months. Figure 25 illustrates this variance across a year by using the daily average values from Figure 22 and Figure 23. The curve in Figure 25 reflects the purpose of HVAC Chillers (i.e. to provide cooling), which is highlighted by the three highest usage levels occurring during the UK summer months of June to August (UK Met Office, 2017b). The lowest three values also align with the UK winter months of December to February. The spring months of March to May show an uneven increase in usage, with March and April only increasing by 2.9% and 12% respectively before May experiences a 46.6% increase. In contrast, the autumn months of September to November demonstrate a more even decrease in usage of 25.4%, 26.1%, and 27.2% respectively.

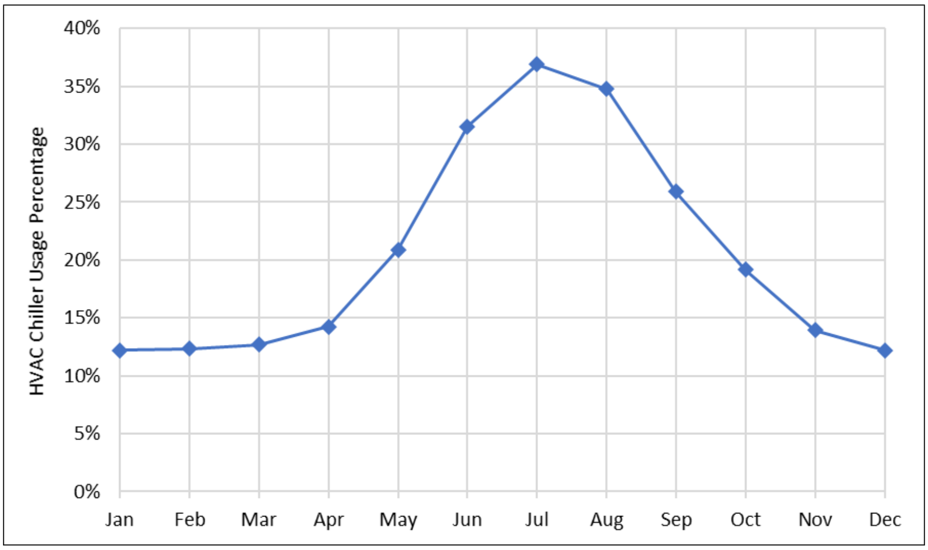

Figure 25 - Daily Average of HVAC Chiller Usage Profile Values per Month

To understand if the variance identified in Figure 25 is linked to outside temperature a correlation was performed as shown in Figure 26. The correlation of average UK temperature against average chiller usage levels resulted in a high R-value of 0.92. This indicates that the usage levels in Figure 25 are influenced by the outside air temperatures, and that the skew towards higher usage in the latter half of year is a result of the chiller meeting the cooling system's prescribed temperature levels during these months.

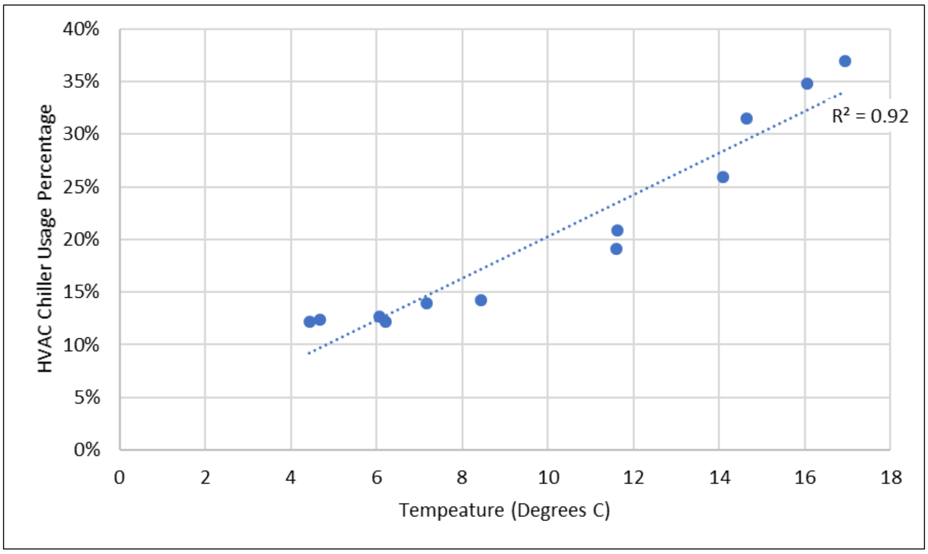

Figure 26 - Correlation of Daily Average of HVAC Chiller Usage Profile Values Against Monthly Average UK Temperatures between 2013 and 2017. Temperature Source: (UK Met Office, 2017c)

The third major feature emerging from the profile method's HVAC chiller profiles is the variance in usage during each day's profile. Figure 27 shows the percentage of time across all twelve profiles that the medium values are above the daily average value. This figure highlights that all twelve profiles follow a consistent usage pattern whereby the chiller usage is lowest in the morning, from midnight to 08:00 (half-hourly period 17), before increasing to its highest usage levels during the remainder of the day from 10:30 to 22:00 (half-hourly periods 22 to 45). Figure 28 shows when the highest and lowest half-hour usage points occur for each profile using the medium values line. Based on Figure 28, the lowest average usage time is 04:00 with a range of between 03:00 and 05:00 (half-hourly periods 7 to 11). While the highest average usage time is 15:30 with a range of between 14:00 to 18:30 (half-hourly periods 29 to 38). To understand how usage variance across the day differs by month Figure 29 shows the minimum and maximum values of the medium usage level. Figure 29 also shows the difference between the minimum and maximum, which highlights that the largest variances across the day of between 11.1 to 14.5 percentage points occur during June to September. The remaining months show a lower variance of between 3.8 to 5.8 percentage points. The four higher variance months also have the highest average usage levels. This indicates that for a hotel chiller the variance between morning and afternoon usage will increase as a greater load is placed on it during these months.

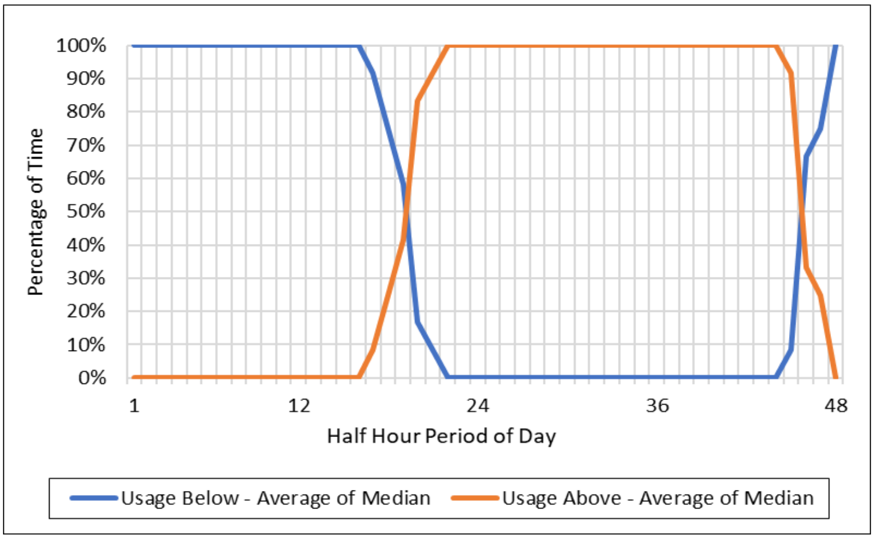

Figure 27 - Percentage of Time Across All Twelve Profiles That the Medium Values Are Above the Daily Average Value

(Base on values from Figure 22 and Figure 23)

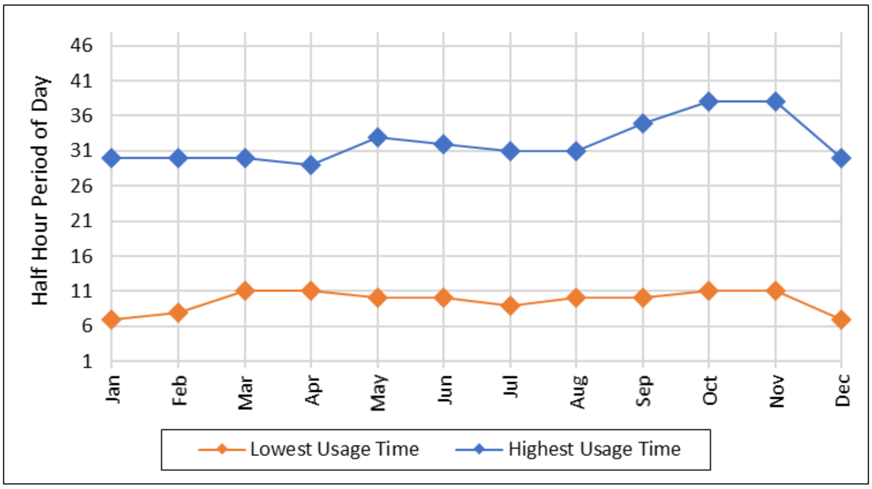

Figure 28 - Highest and Lowest Half-Hourly Usage Time Points by Month

(Base on Median Values from Figure 22 and Figure 23)

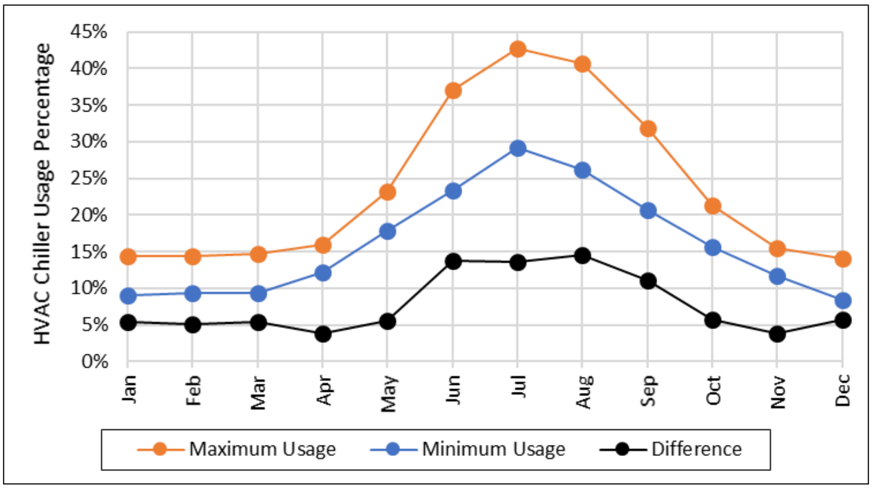

Figure 29 - Medium Usage Level Variance by Month

Validating that the created profiles and the usage trends are representative of UK hotels generally is difficult due to a lack of published data for non-domestic electricity sub-loads. The closest available research is the report by Element Energy (2012) referenced in the chapter introduction. That report generated non-domestic sector electricity usage profiles, separating out the end-usage areas of catering, computing, heating, hot water, HVAC, and lighting to provide indications of DSR potential areas. To understand how the profile method's outcomes compare with other methods, the results shown in Figure 25 will be assessed against the similar outcomes of Element Energy study as shown in Figure 30. This comparison demonstrates that the seasonal trend of the AC & Ventilation end-usage area appearing in Figure 30 is comparable to the hotel chiller profile trends illustrated by Figure 25, with each figure showing low winter usage, autumn having a slightly higher usage level than spring, and summer having the highest usage. However, what is noticeably different between the Element Energy non-domestic electricity sub-loads profile and patterns emerging from the HVAC chiller profiles is when peak usage apparently occurs. In Figure 30 the AC & Ventilation peak occurs at 11:00 during all seasons, as a consequence of Element Energy creating their chart by allocating the overall electricity usage across each end-usage area based on a set percentage usage – i.e. as this is the peak non-domestic usage electricity time, it is also modelled as the assumed peak time for AC & Ventilation load usage too. In contrast Figure 28 shows that the hotel chiller usage in this research actually peaked in the afternoon, at an average time of 15:30. This difference shows the hotel profiles created in this chapter provide greater levels of detail than demonstrated by earlier literature, enabling a more granular understanding of variances in usage across the day and seasons capable of benefitting the wider research field of electricity end-usage undertaken by organisations like Element Energy.

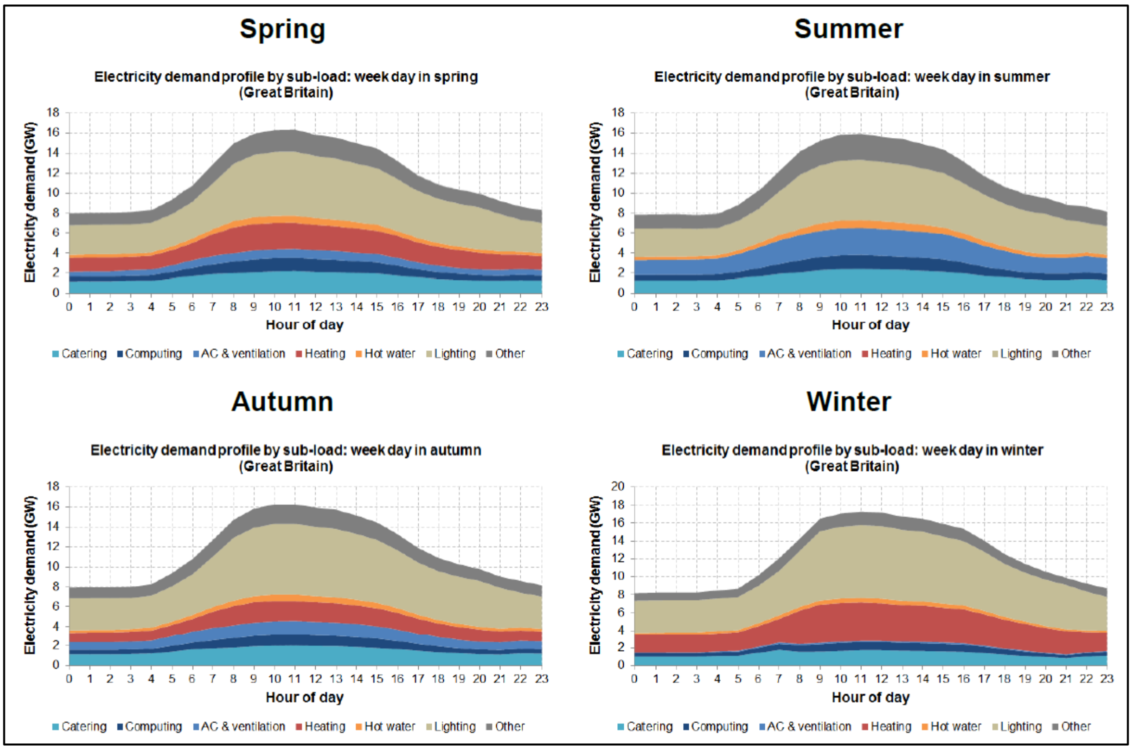

Figure 30 - UK Electricity Commercial Demand Sub-Load Profiles by Season, Source: (Element Energy, 2012)

5.8.2 Reducing DSR Estimation Uncertainty

The purpose of developing the new profile DSR estimation method is to reduce the level of uncertainty during an aggregator's initial desktop assessment of the DSR potential for a new site. One of the key benefits of the profile method is its ability to provide an actual measure for uncertainty at the time of estimation, by reference to the variance between the median, quartile and confidence interval usage values as shown in Figure 22 and Figure 23. This benefit helps overcome a limitation of the existing estimation methods reviewed in chapter 4, as the existing methods each provide a deterministic outcome. This means that uncertainty for the existing methods can only be addressed by retroactively evaluating the DSR potential estimates once actual usage data is made available, i.e. once a site is actually live and usage is capable of being measured. In this section the uncertainty levels of the profiles generated for the HVAC chillers will be examined to assess how they can be utilised during DSR estimation.

Figure 31 provides a box plot of the averaged median, 95% confidence interval and quartile usage levels from Figure 22 and Figure 23 to visually represent the level of uncertainty based on each of the twelve profile's levels of variance (see Table 54 in Appendix E for figure numbers). Figure 32 shows the average percentage points difference between the 95% confidence interval and quartile usage levels. Based on these figures, December to March have the lowest levels of uncertainty with an average variance of 28.7 percentage points between the 95% confidence intervals. In contrast, May and September have the highest uncertainty levels with an average variance of 51.5 percentage points, closely followed by June, July, August, October and November each with an average variance of 45.1 percentage points. These variances mean that the uncertainty of the profile method's estimates will be lower during the winter months when chiller usage is at its lowest level, but higher during the rest of the year, particularly immediately before and after the summer months. Figure 32 also shows that the difference between the upper and lower quartiles is lower and more consistent than the confidence interval variances, with an average variance of 14.3 percentage points across the year and 4.5 points between the highest and lowest levels of variance. This means that while the overall variance in usage levels changes across the year, the variance around the median remains stable.

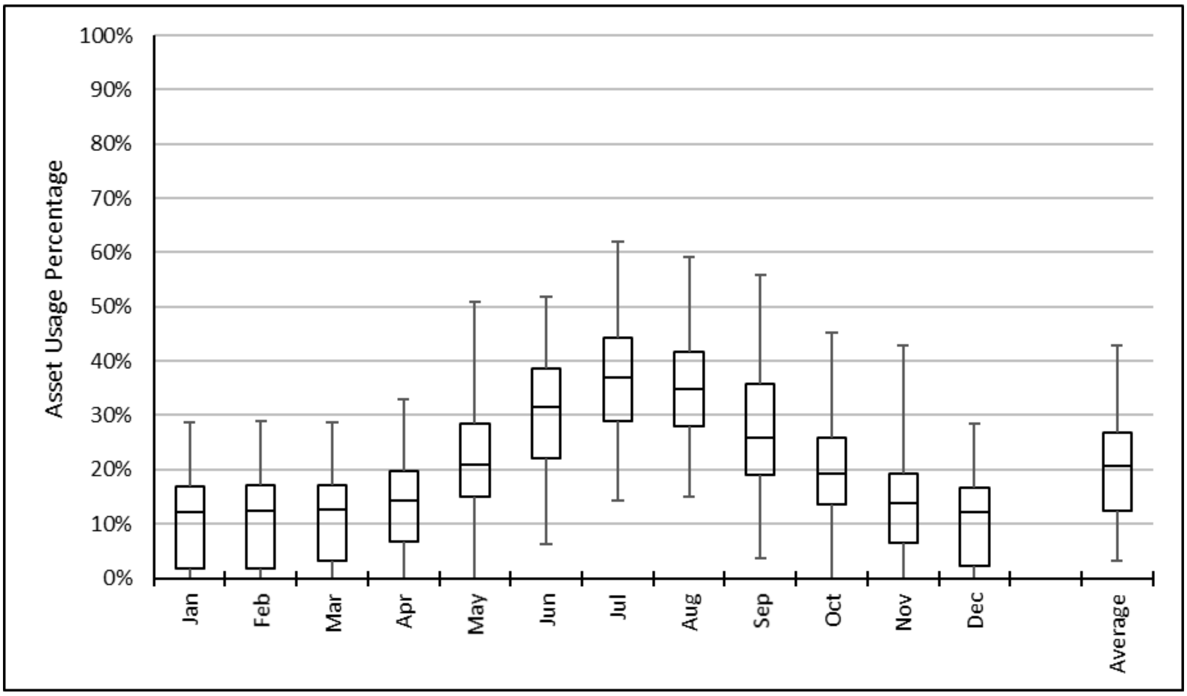

Figure 31 - Boxplot of Hotel HVAC Chiller Percentage Usage

Based on Averages of the Monthly Profiles in Figure 22 and Figure 23. The Box Represents Lower 25%, Median 50%, Upper 75% values, Whiskers Represent Profile Usage at 95% Confidence Interval

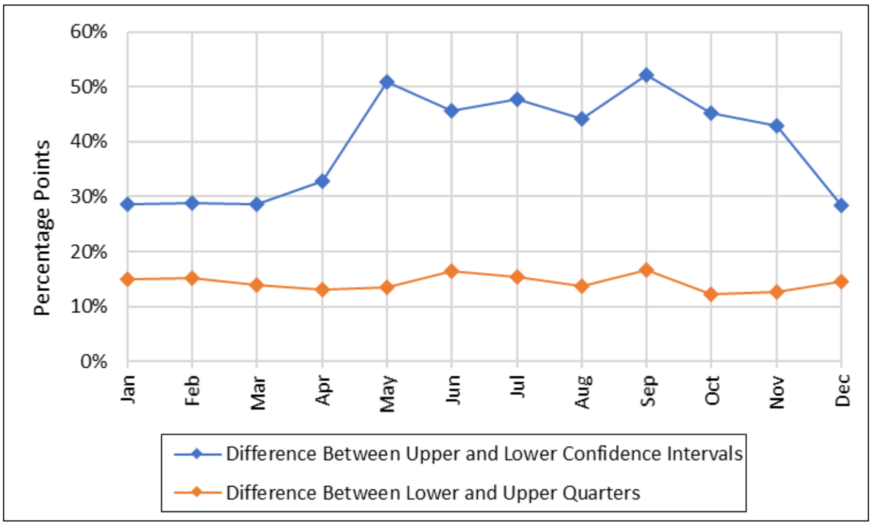

Figure 32 - Average Variance Between Usage Levels of Monthly Profiles in Figure 22 and Figure 23

The variances between the confidence intervals and quartiles help provide an understanding of the uncertainty involved when using the profiles generated by the profile method for estimating a site's DSR potential. This uncertainty will determine the likelihood of the actual usage values being above or below the estimated values. This means that if the user of the profile method is risk adverse, then selecting the lower quartile or confidence interval can mitigate the extent of the uncertainty level, and risk of missing a DSR programme's requirements. For example, using the lower confidence level estimation values (the lower whisker values on Figure 31) will result in the site's actual usage being higher than its estimated usage 97.5% of the time. This allows the user to be confident that there is only a 2.5% chance of not meeting DSR commitments based on the estimate used. This example would be a very risk adverse strategy as it means losing out on the majority of the site's DSR potential due to the low estimation values. Yet it may be that such a risk adverse approach seldom arises, as the profile method is only intended to be used during the early assessment stages for new sites, as outlined in chapter 3. Its intended purpose is to enable an initial understanding of the site's potential for DSR to inform whether further analysis should be undertaken (for example, a site survey) to refine the site's likely actual potential. Therefore, as outlined at the start of section 5.8, it is anticipated that the profile method will likely be used during the initial desktop assessment stage and using the median values, as this provides a balanced DSR potential assessment that can be adjusted up and down through further analysis later if the opportunity is pursued any further.

The variances between the uncertainty levels can also be used to ascertain the likelihood of an estimate being different to the median usage levels. To illustrate this Table 35 demonstrates the results from applying the profile method to estimate the DSR potential of a site with a 200kW chiller. The estimated average kW usage is calculated for each month and usage level by multiplying the percentages used in Figure 31 by 200. As an example, the January median usage percentage is 12.5%. Therefore, the chiller is expected to use an average of 25kW throughout the month of January. Naturally estimates will vary across the day, as per the profiles in Figure 22 and Figure 23. However, for the purposes of supplying a simple example, only the daily average values are used. As the estimates generated for Table 35 use median values, it is expected that for 50% of the time, the actual usage of the chiller will be at least this amount or higher. However, this also means that 50% of the time it could be lower, which represents an area of risk if the estimate is going to be directly used (i.e. not undergo any further analysis), and the site will be used in a DSR programme that penalises for under delivery.

The percentage difference between the kW values at each usage level is calculated as shown in Table 35. This analysis can be used to understand which months will have a higher impact if the selected level of usage is not met. For example, in July the median level estimate shows the chiller usage will be 74kW, yet the lower quartile level shows there is up to a 25% chance that it could be as low as 58kW, which is -22% less than the usage estimate. The lower confidence interval level instead shows that there is up to a 22.5% chance of usage being as low as 29kW, which is -61% less than the median estimate. In comparison the January median usage is 24kW, which drops to 4kW at the lower quartile and then 0kW at the lower confidence interval, comprising a -85% and -100% difference respectively.

Table 35 - Differences Between Usage Levels Example using Profile Method with a 200kW chiller

	Jan	Feb	Mar	Apr	May	Jun	Jul	Aug	Sep	Oct	Nov	Dec
kW usage at Upper CI (2.5% time usage is higher)	57	58	57	66	102	104	124	118	112	90	86	57
Difference between 25% and 2.5% Usage	-57%	-57%	-56%	-57%	-59%	-39%	-41%	-41%	-54%	-58%	-67%	-57%
kW usage at Upper Quartile (25% time usage is higher)	34	34	34	40	57	77	89	83	72	52	38	33
Difference between 50% and 25% Usage	38%	38%	34%	39%	36%	23%	20%	20%	38%	35%	38%	36%
kW usage at Median (50% time usage is higher)	**24**	**25**	**25**	**28**	**42**	**63**	**74**	**70**	**52**	**38**	**28**	**24**
Difference between 50% and 75% Usage	-85%	-85%	-75%	-52%	-29%	-30%	-22%	-19%	-26%	-29%	-53%	-82%
kW usage at Lower Quartile (75% time usage is higher)	4	4	6	14	30	44	58	56	38	27	13	4
Difference between 50% and 97.5% Usage	-100%	-100%	-100%	-100%	-100%	-80%	-61%	-57%	-86%	-100%	-100%	-100%
kW usage at Lower CI (97.5% time usage is higher)	0	0	0	0	0	12	29	30	7	0	0	0

Understanding the difference in kW usage per level enables users of the profile method to have a risk acceptance strategy. For example, the user could decide that they will use the median usage level except where the percentage difference in kW usage between the median and lower quartile is greater than -50%, at which point they will instead use the lower quartile usage level for that month. Based on Table 35, this would result in using the median usage level for months May to October, and the lower quartile usage level for the remaining months. A risk policy decision will depend on many factors necessarily determined by the user (or the user's company), including: (1) what additional analysis will be performed after estimating with the profile method, (2) the penalties for incorrect estimation levels, and (3) an organisation's general risk appetite.

5.8.3 Sensitivity Analysis of Profile Inputs

Section 5.8.2 showed how the profile method enables qualification of an estimate's outcome uncertainty. However, this uncertainty qualification assumed that the created profiles will always use the optimal input values, as determined during the stage 6 optimisation process explained in section 5.6. Therefore, a complete understanding of the potential for uncertainty with the profile method will also require determining the sensitivity of the input values and their impact on the method's outcomes. To enable an assessment of how the inputs that are selected during the profile creation might affect the uncertainty of the profile method's outcomes a one-at-a-time local sensitivity analysis test was undertaken (Saltelli et al., 2008). The sensitivity tests were performed by adjusting each of the input's values as outlined in Table 36, and then re-running the profile creation process to capture the resulting variation in MAPE outcomes.

Table 36 - Summary of Sensitivity Analysis for Profile Method Input Values

Input	Optimal Value Used	Input Value Adjustment	Sensitivity Intervals (including optimal value in **bold**)
How many clusters to use?	3	Adjust number of clusters by +/- 1, 2	1, 2**, 3**, 4, 5
What predictor to use?	Month-of-Year	Adjust for each predictor	Weekend-Weekday, Day-of-Year, Week-of-Year, **Month-of-Year**

Figure 33 summarises the results of the sensitivity tests for each input. To facilitate comparison of sensitivity between inputs, the charts shown in Figure 33 have been normalised by plotting each input variation against the percentage difference in MAPE. These charts show varying sensitivity to inputs within and across the three inputs of site, clusters and predictors. The 'how many clusters to use' input is sensitive when the number of clusters used decreases, as the MAPE value increases by 2.6% when using 2 clusters and 28.9% when using only 1 cluster. In comparison the input is not sensitive when increasing the number of clusters used, as the MAPE changes by only -0.1% when using 4 clusters and -0.2% when using 5 clusters. To assess how these changes impact the profile outcomes, the daily average values are compared (i.e. the overall average of all median usage values from the twelve profiles in Figure 22 and Figure 23). The optimal value of 3 clusters provides the initial baseline, with a daily average of 20.59%. Using 2 clusters causes this average to increase by 0.05 points to 20.64%, whereas using 4 clusters decreases it by -0.33 points to 20.26%. The less than 3% change in MAPE and median averages resulting from changing the number of clusters by +/- 1 indicates that this input has only limited sensitivity to changes around the optimal value.

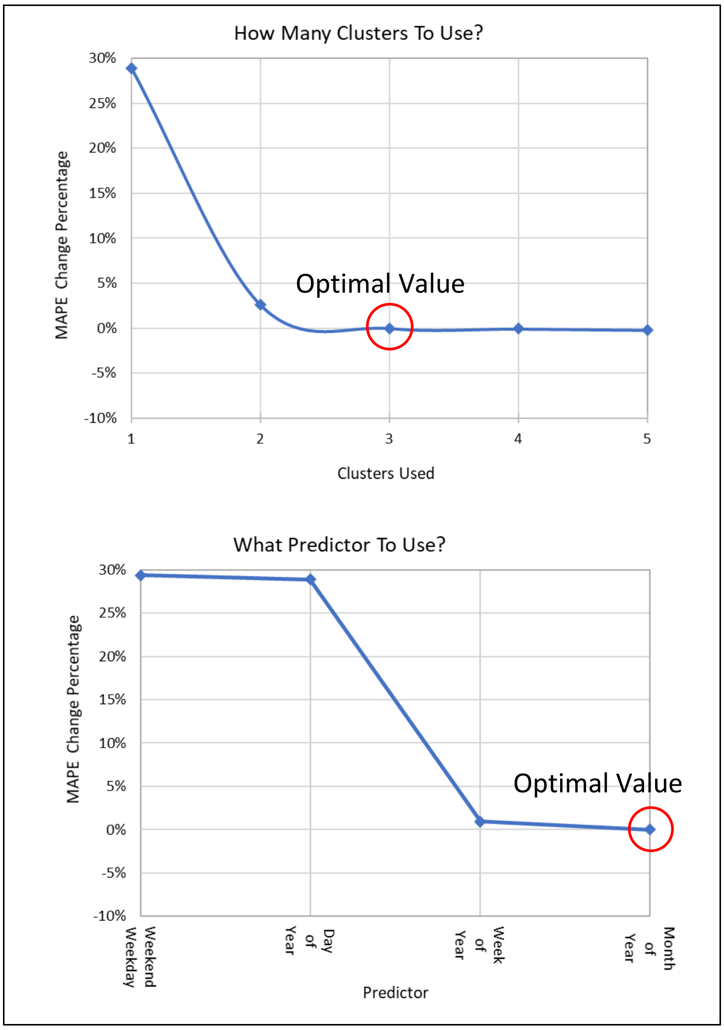

Figure 33 - Profile Method Input Sensitivity Analysis Results

The ‘what predictor to use’ input has four options with a range of MAPE impacts from 0.9% to 29.4%. As seen in Figure 33, there is a distinct change in MAPE impacts between the small 0.9% increase when using the ‘week-of-year’ value compared to the 28.9% increase occurring when using the ‘day-of-week’ value. The different in MAPE values also translate into corresponding impacts on the daily average usage comparison metric. Using the 'week-of-year' value decreases the daily average by -0.13 points to 20.46% whereas using the 'day-of-week' decreases the daily average by -2.12 points to 18.57%. This input has the highest sensitivity of all inputs tested due to the level of MAPE and daily average changes.

The cautionary note about applicability of the sensitivity tests to other profiles applies to all the inputs tested. Each of the two inputs must be chosen using the optimisation process outlined in section 5.6 for each profile creation. As a result, the sensitivity is likely to be different for each profile depending on the data used. For example, the 'what predictor to use' input for HVAC chillers had the largest MAPE impact when changing from using the 'month-of-year' to 'weekend-weekday' value. Yet this is a result of the hotels assessed for this research being identified during the profile creation process as having HVAC usage patterns that vary by month, and not by day, which in turn meant that the weekends and weekdays predictors were deemed unsuitable due to the high error levels caused when tested. If a different predictor value applies (for example, the weekday's predictor had the lowest error levels), then its anticipated that the sensitivity test results would be quite different. Therefore, the sensitivity tests applied for the hotel chiller profiles assist with demonstrating the impact of selecting different values, but should only be utilised for this profile set. If sensitivity tests are undertaken on future profile creations, then these results could be used to determine whether there are common sensitivity patterns between profiles for different business categories and assets. Putting aside the potential for further research, the sensitivity analysis undertaken for this thesis helps validate the optimisation process as the sensitivity results show that the values closest to the selected optimal values for each input have limited impacts on the MAPE outcomes.

5.8.4 Evaluation of New Method Against Existing Estimation Methods

This final section looks at how the new profile method compares to the four existing DSR estimation methods examined in chapter 4. To understand how the new profile method compares with existing methods, the same comparison technique from chapter 4 has been applied for the twelve HVAC chiller profiles described in section 5.7.1. The results of this comparison are evaluated in three ways: (1) comparing the MAPE and MBE outcomes, (2) comparing the method costs against error levels, and (3) reviewing the estimation patterns produced by all methods.

The first evaluation of the new profile method against the existing methods compares the MAPE and MBE results of DSR estimations for two years for both test hotels (Regents Park and Bristol

Royal). To complete this test for the profile method, the hotel site being estimated was removed from the profile creation process to enable an estimate for an unknown hotel using the information about chiller usage for known hotels. As discussed in section 5.8.2, a primary benefit of the profile method is the ability to determine the level of acceptable uncertainty during estimation applying different usage level values (i.e. median, upper and lower quartile and 95% confidence interval). To compare the profile method's uncertainty levels with the four existing DSR estimation methods the previously used Figure 13 in chapter 4 has been used for Figure 34 and then extended to include the average, minimum, and maximum MAPE and MBE results for each profile usage level (see Table 55 in Appendix E for figure numbers).

To understand the impact of the five different usage levels generated for each profile they were each tested which resulted in a range of MAPE and MBE values, as shown in Figure 34. Looking at the median usage level first, this generated the lowest MAPE and MBE values at 46.5% and 1.1% respectively. When compared to the existing methods MAPEs, the profile method's result is the second lowest behind M3 at 38.8%. The positive MBE value indicates that using the median usage level will likely result in a small amount of overestimation, which contrasts with the existing methods which all underestimate except for M1-V1. While the median usage level provides a balanced estimate, the profile method also allows for other estimation levels to be utilised based on the user's level of acceptable uncertainty as discussed in section 5.8.2. The impact of using the different usage level values is illustrated in Figure 34, which shows the four non-median usage levels causing the MAPE to increase. The reason for the MAPE change is highlighted by the MBE values, which follow a pattern of increasing (greater positive bias) when using the upper usage levels, and decreasing (greater negative bias) when using the lower usage levels. While the MBE and MAPE values are larger when not using the median usage level, this is not necessarily a negative outcome as the larger values highlight how estimation bias can be used to measure uncertainty when selecting which usage level to use. If the user selects the lower confidence interval usage level for the estimation (and so ensures that 97.5% of the time the actual usage will meet or exceed the estimate), then the MBE value of -83.6% obtained during this evaluation helps reinforce that the estimations were consistently under the actual usage levels.

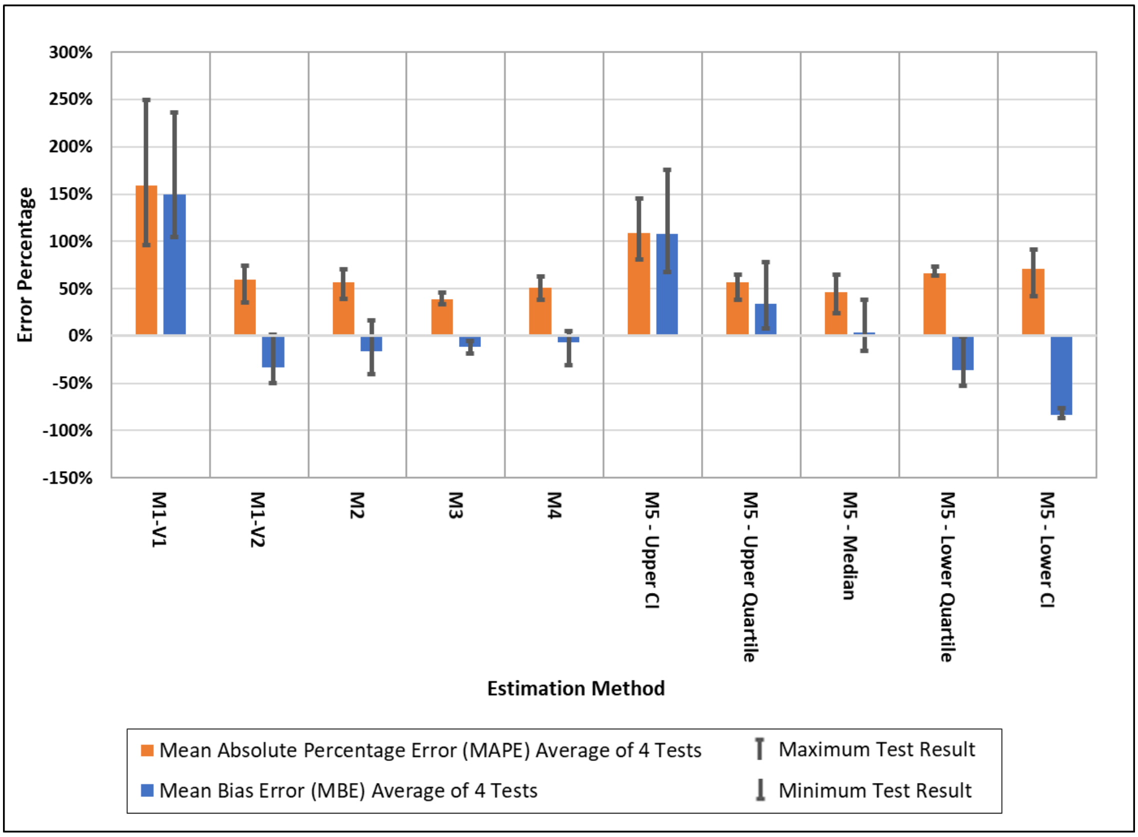

Figure 34 - Summary of Existing (M1-M4) and New (M5) DSR Estimation Methods Based on Average MAPE and MBE Results with Minimum and Maximum Bars from 4 Years of Data From 2 Hotels

Abbreviation key:

M1-V1 = Method 1- Variation 1 - Minimum information using set percentage of asset usage
M1-V2 = Method 1- Variation 2 - Utilise baseload calculation with set usage percentage
M2 = Method 2 - Baseline comparison using cluster analysis
M3 = Method 3 - Regression analysis utilising historical DSR event outcomes
M4 = Method 4 - Building energy modelling
M5 = Method 5 - Usage Profiles

The second evaluation compares the cost of running the new profile method against the existing methods' costs and expected levels of estimation error. As outlined in section 4.2.3, this comparison helps provide context for usage of the methods when balancing cost against acceptable error levels. This is an important factor for DSR aggregators, as reviewed in chapter 3, who are performing these estimations daily with costs primarily covering employee time as shown in Table 21 and Table 34. Figure 35 extends the cost against error appearing at Figure 15 from section 4.2.3 by adding the M5 profile method and using for this method the MAPE values from the median usage levels, and costs from Table 34. The chart shows that the new profile method is second cheapest with a per usage cost of £26, which is £16 more expensive than M1-V1 and £4 cheaper than M1-V2 and M2. As the profile method also has the second lowest MAPE value of 46% which is 7 points higher than the M3 method that has the lowest MAPE of 39%.

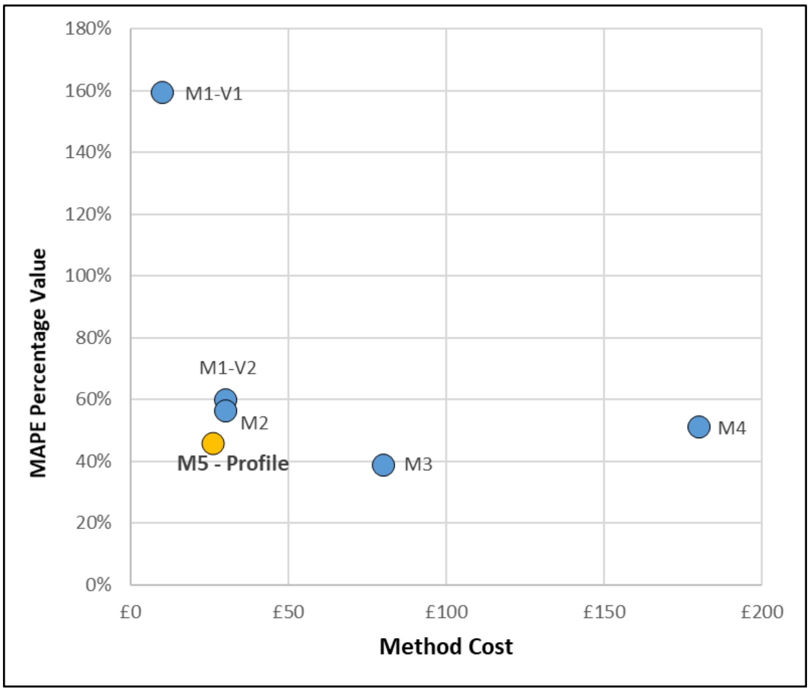

Figure 35 - Comparison of Estimation Method Error versus Usage Cost
Based on chapter 4 - Figure 15 with the addition of the Profile Method as developed in this chapter.

The third evaluation compares how the estimation outcomes for each method differ when viewed against actual usage. The Bristol Royal Hotel's 2013 dataset as shown in Figure 36 was used for this evaluation. Only one dataset is used for the comparison in order to clearly illustrate any differences between the estimation methods. The figure shows that there is an overall trend of underestimating, which reflects how all methods have a negative MBE except for M1-V1. Figure 36 helps explain the cause of this, as it can be seen from this comparison that most of the estimations track around the actual usage level for all months except June and July. July is expressly different, as all estimates except M1-V1 are approximately a third lower than actual usage. Figure 36 will be used to describe how each method's estimation varies across the year.

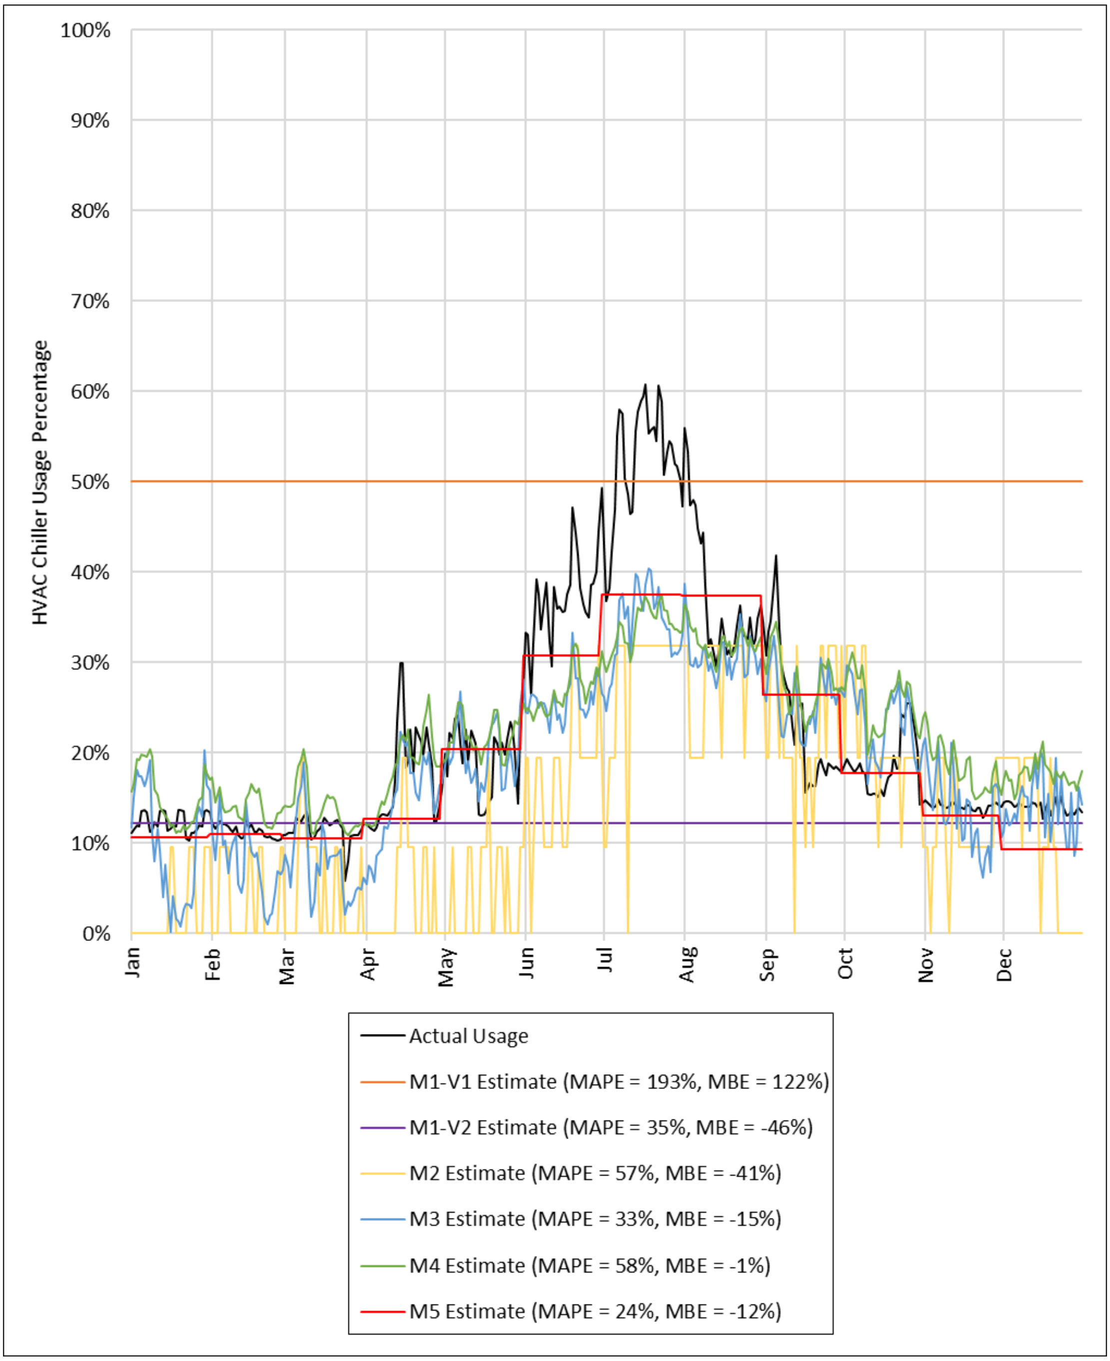

Figure 36 - Comparison of Actual Usage against Estimations Methods Averaged by Day using results from the Bristol Royal Hotel 2013 HVAC Chiller Dataset

M1-V1, described in section 4.1.2.1, is based on assuming a fixed 50% usage level of the assessed asset's maximum kW rating. Due to this simple calculation, M1-V1 has the lowest per use cost of £10. However, it has the highest MAPE error of 193%, and an MBE of 122%, as visually represented in Figure 36 by the overestimation during all months except June and July. M1-V2 is a variation of M1-V1 and performs a pre-calculation step by first determining the site's baseload usage, which is calculated at a 5% percentile of overall electricity usage levels for a year. A percentage (default 10%) of the baseload is then deemed to be used by the DSR asset. This approach reduced the MAPE to 35%, the third lowest error level of the methods for Bristol Royal Hotel in 2013, and produced the third equally lowest per use cost of £30 (matching M2's estimated usage costs). The baseline approach still results in a consistent estimation level, which matches M1-V1 as seen in Figure 36, and results in a -46% MBE due to underpredictions during the summer months. However, as the estimation closely aligns to the non-summer months, M1-V1 has a significantly lower MAPE error than M1-V2.

M2, described in section 4.1.2.2, used clustering to identify when the chiller was operating based on the site's overall electricity usage records. As this method assumed that the lowest usage cluster represents when the chiller is not operating, the estimation levels, as set out in Figure 36, show a distinctive up and down pattern, which is due to the days that are deemed to have no usage. This method underestimated usage with a -41% MBE and produced the third-highest MAPE of 57% and the third equally lowest per use cost of £30 (matching M1-V2's usage costs).

M3, described in section 4.1.2.3, uses regression analysis of past DSR event outcomes in conjunction with outside air temperature to predict the expected level of chiller usage over a year. The estimation line in Figure 36 for this method varies across the year due to its usage of the outside air temperature to estimate chiller usage. This method had the second lowest MAPE of 33% for the Bristol Royal Hotel in 2013 and an overall underestimation bias based on an MBE of -15%. Although this method generated the second lowest error level, it does have the second highest usage cost of £80 due to informational requirements and time needed to perform the calculations.

M4, described in section 4.1.2.4, used an Energy Building Model to simulate the expected kW usage level of the chiller for the site. The energy model used the same outside air temperature dataset as M3, as seen from Figure 36, and the estimations from these two methods also follow a similar trend from May to October. Outside of these months the simulation estimate is generally higher than the actual usage levels, which results in this method receiving the second highest MAPE of 58% with an MBE of -1%. This method also has the highest usage cost of £180 as a result of the time required to create the energy model.

The M5 profile method outlined in this chapter provided the lowest error level for the Bristol Royal Hotel in 2013, with a MAPE of 27% and a MBE of -12% when using the median values from the HVAC chiller profiles. As the profiles created for hotel chillers each cover a month of the year, the estimation trend line in Figure 36 does not have the same intra-day variances as seen in M2-M4. If the profile optimisation process had instead selected the day-of-week predictor for example, then this would have resulted in greater variations. Generating the second lowest cost to run at £26 and the lowest MAPE for this site in 2013 helps validate the profile method as a suitable new method for DSR estimation.

5.9 Chapter Conclusion

This chapter sought to develop and evaluate a new DSR estimation method with capability for reducing the levels of uncertainty in the resulting outputs. The resulting new method was developed by creating a set of load profiles using detailed usage information (e.g. from sub-meters), representative of common usage patterns for electrical assets, which for the purpose of this research was hotel HVAC chillers. These profiles are intended for use in then determining the likely usage levels for similar assets at sites where detailed usage information is not readily available. Choosing to resolve this chapter's objective by using profiles to create the 'profile method' for DSR estimation during initial assessment phases for new sites was informed by publishing research, as reviewed in section 2.4 on building energy usage estimation approaches. While there has been extensive research on using load profiles for different applications, prior to this thesis, there has not been any direct research into the capability for using load profiles in enabling new methods for site level DSR potential estimation. The closest usage of profiles for DSR was by Element Energy (2012) who used country-level profiles of major end-usage areas to calculate the DSR potential for the UK. Therefore, this chapter made use of that existing load profile research to develop and assess a new method of site-level DSR estimation informed by usage profiles. This conclusion evaluates the research outcomes from this chapter in two sections. Section 5.9.1 reviews each stage of the profile method development process, summarising the major outcomes and findings. Section 5.9.3 provides a summary of the overall findings of this chapter to highlight the implications of this research.

5.9.1 Profile Method Development Process Review

A seven-stage process as outlined in Figure 16 was undertaken to develop and then assess the profile method, as a new approach for estimating a new site's DSR potential. The key outcomes of each stage are summarised as follows:

Stage 1, section 5.1 outlined the data sources used for developing the profiles. This research developed proof of concept usage profiles using sub-metered data from five UK hotel HVAC chillers, as obtained from the research partner KiWi Power. Across the five hotel sites 4,012 days of data was gathered, consisting of minute interval kW usage readings across each day.

Stage 2, section 5.2 addressed data preparation, including converting and normalising the data, and resolving the treatment and management of missing values. From this process, the minute interval kW chiller usage readings were first converted into half-hourly kW usage intervals before being normalised as a percentage of the maximum usage kW usage rating of the chiller. The data was then analysed to determine the level of missing values. This analysis identified that 63 of the 4,012 days (1.57%) of data contained between 1 to 25 missing half-hourly usage values. This raised two method development questions: firstly, what error is introduced by using interpolation to fix missing half-hourly gaps, and secondly, what is the overall impact on the profile outcomes of fixing the identified gaps in the data.

To answer the first question an assessment was undertaken to determine the level of error that would be experienced when fixing different scales of missing values using interpolation. The assessment was performed by purposefully making gaps of between 1 to 48 values in a half-hourly hotel chiller dataset. These gaps were then fixed using interpolation, and evaluated to understand the level of error between the fixed gap and actual data that had been removed. The results showed a gradually increasing error level as the gap increased, started at 6.7% when fixing one missing value, and increasing to 22.1% when fixing 48 sequentially missing values. The interpolation method was selected in reliance on the research by Sluiter (2009) and Meijering (2002) which deals with addressing missing data in time series. Applying these findings in this research helped expand knowledge about the applicability of the interpolation method for electricity usage datasets with missing values.

Addressing the second question (consequences on outcomes of fixing data gaps) required looking in detail at the impact of using interpolation on the profile outcomes. This question was addressed as described in section 5.6.2.1 of stage six by assessing the outcomes arising from applying different levels of data fixes. The results showed that fixing half-hour data gaps of between 1 to 6 produced the lowest MAPE of 58.96% while not fixing any gaps generated the highest MAPE of 59.25%. Based on these results it is recommended for future profiles that data gaps of this size are fixed.

Stage 3, section 5.3 reviewed options for creating training and testing datasets for the purposes of profile creation. Evaluating these profiles involved splitting the chiller usage data into a training dataset for profile creation, and a testing dataset for evaluation. A method development question was raised at this stage regarding which approach should be used for the purposes of this research based on different options being proposed in existing literature. Therefore three methods were selected for assessment to determine which one should be used: (1) the traditional 'Out-Of-

Sample' approach that splits the dataset based by using a percentage of the latest data in the time series as the testing dataset, (2) the 'K-fold Cross Validation Using Sites' approach whereby one site is used as the testing dataset and the remaining sites form the training dataset, a process which is then repeated until all sites have been used for testing, and (3) the 'K-fold Cross Validation Using Random Selection' approach whereby each day of data is assigned a random value between 1 to 4 before creating the testing dataset using all days assigned to 1 and then the training dataset using days 2 to 4, a process which is then repeated until all random values have been used for the testing dataset. Section 5.6.2.2 of stage 6 assessed each option, and determined that option 3 - 'K-fold Cross Validation Using Random Selection' would be utilised as the preferred data selection approach for creation of profiles for DSR estimation.

Stage 4, section 5.4 outlined the profile creation process, which included reviewing clustering methods, deciding how the predictors for usage profiles were selected, creation of weighted distributions for profile creation, and managing non-normal distributions. Creating profiles involved first grouping the available data into similar patterns using the K-means technique based on it being a proven method for time-series clustering that continues to perform well, is understandable, efficient and scalable. A key input to the K-means method is determining the number of clusters to use. This raised the first recurring question, which will need to be addressed each time a new profile is created. For the purposes of this research, this question was assessed and resolved in section 5.6.3.2 of stage 6, resulting in the selection of three clusters for the hotel chiller profiles.

Creating usable profiles then required selection of an appropriate predictor, to create the link between the profile and the subject that it is being applied to. The new profile method for DSR estimation has two default predictors in the form of the business category (i.e. hotel) and asset (i.e. chiller) due to the data used to create the usage profiles. However, a third predictor is required to determine how the profiles created for each business and asset type should be applied. As the application of the profiles for DSR estimation must be done with minimum information during the initial phases for assessing a new site, only a date based input was available for the third predictor. This raised the second recurring question: the need to determine the optimal predictor for each profile creation. For this research, this involved assessing as four options Week-of-Year, Month-of-Year, Day-of-Year, Weekend-Weekday. The assessment described in section 5.6.3.3 of stage 6 identified that the Month-of-Year predictor provided the lowest overall error level for the hotel chillers.

As each predictor has a different number of profiles (e.g. Week-of-Year = 52 profiles, Month-of-Year = 12 profiles), each profile required the creation of a specific distribution of usage values based on the clustering outcome. These distributions were created using stochastic process and involve selecting random values from the clusters associated with each predictor value. A method development question was raised at this stage on the size of the random sample. This question was assessed in section 5.6.2.3 of stage 6 resulting in the selection of 3000 samples based on the next larger size having margin benefit.

As the distributions were found to be non-normal and primarily positively skewed, using mean and standard deviation values could not present accurate outcomes for the load profile creation. Instead percentile values were used as a solution, as recommended by Laerd (2017) who explains that these provide a more realistic measure when using non-normal distributions. To assess variances within each profile, five usage levels were determined using percentile quartiles (25%, 50%, 75%) and 95% confidence intervals values (2.5% and 97.5%). Using these measures, the profiles for each month of the year were created as shown in Figure 22 and Figure 23.

Stage 5, section 5.5 evaluated the created profiles using the same methods described and applied for chapter 4, namely the MAPE and MBE. The MAPE value provides an absolute level of error, while the MBE value will show if there is a positive or negative error bias. This evaluation was undertaken by using the profiles to create estimates, which were then compared to actual values in the testing dataset using the MAPE and MBE methods. The evaluation process was used during the stage 6 optimisation and stage 7 to produce the final results, as reviewed in section 5.8.

Stage 6, section 5.6 performed an optimisation process to address the three-method development, and two recurring input value selection questions were raised during stages 1-5. The optimisation process was undertaken by selecting the values to be tested for each question, and then creating and evaluating profiles that covered all possible input value combinations. This created 9,120 combinations and took approximately 22 hours to process. The optimisation outcomes were used to address the five questions as previously reviewed in this section. A summary of questions and answers is provided in Table 37.

Table 37 - Summary of Questions Addressed During Stage 6 Optimisation

Question	Answer
Which missing data option should be used?	Use the ‘fix single gaps’ option
Which training and testing data selection method should be used?	Use the ‘K-fold Cross Validation Random Selection’ method
What stochastic sample size to use?	Use a sample size of 3000
How many clusters to use?	For this profile used 3 clusters
What predictor to use?	For this profile use Month-of-Year predictor

Stage 7, section 5.7 generated DSR estimations using the profile method to enable it to be compared against four existing DSR methods that were reviewed in chapter 4. DSR estimations were created using the new profile method for the two test hotels (i.e. Bristol Royal and Regents Park) assessed in chapter 4 by excluding one hotel at a time from the profile creation process. Additionally, the cost of running the new profile method was calculated at £26 per use, which included both the time to run the method and also maintain and create new profiles.

5.9.2 Potential Applicability of Profile Method to Different Asset and Building Types

The DSR profile estimation method has been created and tested using data from five UK hotel HVAC chillers. This raises as a question how applicable this method may be for other asset and building types. Assessing the applicability of this method to different asset and building types is beyond the scope of this research due to time and data limitations. However, it is possible to undertake a high-level theoretical review to understand the DSR profile estimation method’s potential for wider applicability. This section undertakes this review through three examples that examine use of the method for a different asset and building type combination to determine if the profiling method could be applicable for estimating DSR potential of other assets. The asset types used are based on the three non-domestic electricity usage areas identified in section 2.1.3 as showing the highest potential for providing DSR feasibility: hot water tanks, HVAC, and pumps and compressors. The building types selected for each asset example are based on previous research of their usage for DSR and consist of a domestic building for hot water tanks (Jack et al., 2018), an office building for HVAC (Chen et al., 2017), and a water pumping station for pumps and compressors (Menke et al., 2016).

Hot water tanks were identified at section 2.1.3 as providing a high level of flexibility for DSR due to their thermal store, which allows temporary interruption of water heating to occur without impacting users. Research by (Jack et al., 2018) on the flexibility offered by domestic hot water tanks for DSR showed that usage patterns across a day would vary between being on for many short periods (less than 15 minutes) and a few long periods (between 1 to 2 hours). The long periods resulted through high usage of hot water (for example, having a bath or washing the dishes) that required new cold water to be heated. The short periods arose from either smaller usage of hot water (for example, hand washing) or through thermal loss of the stored hot water that resulted in the temperature dropping below a set point. When examining usage over a year and across multiple buildings Jack et al. (2018) identified consistent patterns of high usage between 07:00 to 08:00 and 17:00 to 19:00, medium usage between 08:00 to 17:00, and low overnight usage. These consistent use patterns mean that the DSR profile estimation method should be able to determine the DSR potential of hot water tanks in domestic buildings. Potential prediction parameters for usage of the profiles could include weekday vs weekend, building occupancy (for example, family vs non-family), and building type (for example, standalone vs apartment).

The second building and asset combination review looks at HVAC usage in office buildings. The HVAC profiles created in this chapter showed that the cooling element of hotel HVAC systems followed a monthly pattern, with limited differences between weekday and weekends. The lack of variance between weekday and weekends reflects how the hotels in this research were open all week for business. In contrast, research by Chen et al. (2017) on HVAC in 4 office buildings showed a different usage pattern whereby there was consistent high usage during weekdays and low usage during the weekends. However, this still had consistent patterns of usage which means that the DSR profile estimation should be suitable for HVAC in office buildings. It is also likely that the profile predictors can be based on a mix of weekend vs weekday and week or month parameters informed by the office buildings' defined usage patterns.

The final asset and building type combination to be review is a water pump at a water pumping station. As reviewed in section 2.3, research by Menke et al. (2016) showed that water pumping stations were found to have a medium to high flexibility level as these are normally used to move large amounts of water into storage reserves, which enables temporary interruption with minimal impact and makes them suitable for DSR. Their research examined a site that consisted of two main pumps (each rated at 178 kW), one of which would normally be operational at different levels of load for the majority of the day and the second being periodically used when the first one could not meet demand. These usage patterns indicate that the DSR profile estimation method could be successfully used to predict the DSR potential of the pumps if other sites follow a similar usage pattern to the site researched by Menke et al. (2016). While their research did not provide indications of variance across the year, if it is assumed that usage of water pumps at pumping stations is relatively consistent, then the profiles created from this data could potentially have predictors based on the day of the week, generating 7 profiles.

Application of the DSR profile method in these three examples indicates that it is flexible enough to manage different scenarios. Validation would require access to sufficient data to create and test the profiles. However, this theoretical review informed by others' research shows how it is likely to be applicable to different building types, assets and operational patterns.

5.9.3 Summary of Profile Method Development Outcomes

In section 5.8 the results of the profile development method were reviewed and discussed starting with an overview of the actual profiles created. The overview identified three key features of the hotel chiller profiles: (1) a 'rebound' effect, which caused spiking of usage during winter in the morning and evening at the Maida Vale hotel. Data analysis identified that these spikes resulted from the hotel chiller temporarily working higher than normal as it returned cooling conditions to normal after having been turned off at night and during Triad periods. (2) the hotel chiller systems have an expected behaviour pattern of highest usage during summer and lowest usage during winter. The link between chiller usage and weather conditions was validated through a correlation of average UK outdoor temperatures against chiller usage, which produced a high R-value of 0.92. (3) all twelve profiles follow a consistent usage pattern whereby the chiller usage is lowest in the morning from midnight to 08:00 before increasing to the highest usage levels from 10:30 to 22:00. The lowest average usage time is 04:00, while the highest average usage time is 15:30. The variance between highest and lowest usage times was greatest from June to September with a difference of 11.1 to 14.5 percentage points while the remaining months show a lower variance of between 3.8 to 5.8 percentage points.

These features have several implications for DSR. Where DSR programmes have inflexible reduction commitments like STOR (see section 2.1.1), the variations across the day and months will impact the level of perceived DSR potential. The profile variation across a day means that if the DSR programme is like STOR, and only allows a single daily kW reduction amount, then the lowest usage level will need to be used in completing estimates. The variation across months will also impact the amount that can be reduced, depending on how inflexible the DSR programme is.

If the programme allows only one fixed reduction amount across the year, and has strict penalties for missing it, then this will restrict the DSR estimation towards only using only the lowest winter profiles. If the DSR programme has limited flexibility, like the STOR programme's splitting of the year into six separately tendered for periods then the different usage levels in the profiles can be used to increase estimation by having different reduction amounts for each period. Where DSR programmes involve higher levels of flexibility, like FFR (see section 2.1.1), then usage profiles could be used to improve the DSR estimation potential, by varying the reduction targets across the day and month.

In addition to providing important insights into usage patterns across the day and over months the profile method's usage profiles also enable uncertainty arising during the estimation process to be considered in making decisions. The profiles' median, quartile and confidence interval usage values enable the level of potential variance in the estimation to be ascertained, which can then be utilised to determine how certain the estimated DSR potential of a site is likely to be, and what risks might arise in proceeding further. The implications of the capability for assessing uncertainty on the potential for DSR estimation arises in the ability of the user to determine an acceptable risk strategy, informed by anticipated percentage variances between the profile usage levels. As demonstrated in section 5.8.2 the user (or users company) can select a level of acceptable risk based on percentage likelihood of the actual usage value being lower than the estimate. This ability to adjust the estimate based on the level of uncertainty provides greater flexibility than the existing deterministic DSR estimations methods reviewed in chapter 4.

To understand how the new profile method performed against the four existing estimation methods the assessment approach used in chapter 4 is applied using the new profiles generated, and the results from all methods then compared. The comparison showed that DSR estimations using the new profiles' median usage values resulted in the second lowest MAPE of 46.5% and second lowest per usage cost of £26. This implies that when using the profile method in a basic deterministic form (e.g. only using the median level usage values) it is suitable as an approach for DSR estimation, and one that outperforms three other methods.

Figure 34 also showed the additional benefits of the new profile method when using each of the upper and lower quartile and confidence interval profile usage levels for DSR estimation, which resulted in varying levels of MBE outcomes. As the MBE value provides a measure of under or overestimation, the different profile usage levels can be selected based on the intended outcome from the estimate. If the estimate needs to be conservative, for example, then selecting the lower quartile will forecast usage values that will be under actual real usage values 75% of the time. Alternatively, if there is a desire to understand a site’s highest potential usage levels, then selecting the upper confidence interval values will show the usage values that occur less than 2.5% of the time. This flexibility in conjunction with the ability to manage the acceptable level of uncertainty by using the different profile usage enables users of the profile method to have greater control of their DSR estimation.

This chapter’s development of a new profile-based DSR estimation method has three key implications. Firstly, using load profiles for DSR estimation in its basic deterministic form has been proven to have the second lowest error level of the five methods tested, and the second lowest cost. Combining these two metrics results in the profile method having the best overall balance of error and cost as shown in Figure 34, justifying its recommendation for future usage for estimating a site’s DSR potential. Secondly, the profile method provides the ability to manage uncertainty at the time of estimation, which distinguishes it from existing methods. Reducing the level of estimation uncertainty could help improve decision making when determining the suitability of a new site. Thirdly, the profile method has been proven using hotel chiller usage data only. Further validating this method and its outcomes will depend on creating through this method new profiles for different business categories and electrical assets.

6 Conclusion

This EngD research thesis has investigated and improved the understanding of DSR aggregators, DSR estimation methods, and developed a new profile-based estimation method. The primary motivation for this thesis is to improve the understanding of the impact and effectiveness of different DSR estimation methods on assessing the suitability of new sites to help increase usage of DSR, which in turn will play a role in helping to achieve a low carbon future. This future role for DSR and its capability for enabling a low carbon outlook is reflected in an SEDC (2015) report which categorises DSR as a requirement for meeting the EU 2030 Energy Strategy objective of renewable energy achieving at least a 27% share of overall energy consumption. In the UK, the National Grid considers DSR as a key enabler for managing future grid variability, which is reflected in their Power Response programme and its aims of achieving 30-50% of network balancing capability from DSR by 2020 (National Grid, 2017).

DSR aggregators play a key role in currently providing over 80% of DSR capacity, a role which is likely to continue (Ofgem, 2017a; SEDC, 2017). Aggregators act as intermediaries by providing services to end users who would otherwise be unable to participate in DSR due to DSR programme requirements (for example, minimum kW reduction, complex control systems, bidding for capacity). As an intermediary in the process, aggregators need to ensure that an end user's site is suitable for DSR, otherwise there might be financial and reputational impacts for both parties. A crucial element in determining suitability is understanding the potential financial returns. Gaining this understanding requires estimating the site's potential for DSR, which is informed by the available electricity assets capable of being used for DSR.

A review of existing literature in chapter 2 on how aggregators determine the suitability of a new site and the DSR estimation methods used revealed very limited research in this area. Therefore, through three research objectives, this thesis has focused on increasing the understanding about how an aggregator determines site suitability, what existing DSR estimation methods are available and how these compare, and then developing a new load profile-based DSR estimation method to reduce uncertainty and improve new site DSR suitability assessments. This concluding chapter provides a review of each of the research objectives in section 6.1 before focusing on the overall implications of the research in section 6.2, and finishing in section 6.3 with an overview of the research limitations and future potential research areas.

6.1 Research Objective Review

This section provides a brief overview of each of the three research objectives, covering the objective's purpose, what research was undertaken, and a summary of the findings.

6.1.1 Objective 1 Review - DSR Aggregators: How They Decide Customer Suitability

Chapter 3 addressed the first research objective *'To map out the criteria used by an aggregator to determine site suitability for DSR'*. The purpose of this objective was to understand the current site suitability assessment process in order to provide a knowledge foundation for addressing this research project's overall aim of improving DSR uptake. As DSR aggregators are the primary conduit for DSR, they are an appropriate area to focus on for the purposes of understanding how new sites are assessed for DSR potential. However, there is very limited existing literature on aggregators as reviewed in section 2.2. None of that research looks at how aggregators decide if a site is suitable for DSR. Therefore, this objective was required to address this knowledge gap, and provide a foundation for the remaining two objectives.

To enable an adequate knowledge foundation, two streams of investigation were undertaken. Firstly, a detailed analysis of KiWi Power's new site assessment workflow system to gain an in-depth understanding of this UK aggregator's assessment process, client base, and the reasons why sites are, or are not, categorised as suitable for DSR. Secondly, based on the workflow system review, semi-structured interviews were undertaken with twelve KiWi Power employees to gain further insights into each stage of the workflow process, and the reasons behind site suitability assessment outcomes.

The results of the first objective help address the identified knowledge gap about how DSR aggregators assess suitability with the following key findings: (1) only 36% of new sites complete the process and are enabled for DSR; (2) the decision to not progress with the assessment is made by the client 72% of the time; (3) the primary reason for sites being lost is due to 'loss of interest' by the client; (4) a turndown site would require at least 200kW of turndown potential, from a minimum of two assets each providing at least 100kW; (5) a generator-based DSR site would require at least 500kW; (6) understanding the DSR potential of a site's assets is the highest priority during the assessment process.

6.1.2 Objective 2 Review - A Comparative Analysis of Demand Side Response Estimation Methods

Chapter 4 addressed the second research objective '*To perform a comparison of the outcome uncertainty in DSR potential estimation methods, evaluated against the level of informational requirements of those methods*'. The purpose of this objective was to understand how four known non-domestic DSR estimation methods compared in relation to uncertainty levels based on their input requirements. Section 2.3 identified literature for four existing DSR estimation methods. However, no previous research was found that compared all these methods together. Therefore, this objective involved undertaking comparisons because addressing the overall aim of improving the DSR estimation process necessitated first knowing and understanding the existing methods available, and how they compared.

The comparison was performed by using each of the four methods to estimate the DSR potential of HVAC chiller assets at two hotels over two years. The chiller asset type was selected due to the high potential these assets offer for DSR, as outlined in section 2.1.3. The hotels and years were selected due to KiWi Power being able to provide minute interval kW usage data of the chillers from sub-meters they installed at these sites in 2012. Having actual usage data meant that the estimations could be evaluated using the MAPE and MBE methods to assess the level of prediction error. Additionally, the information input requirements of each method were assessed to determine a per usage cost to enable evaluation of each method's error level against the cost to use. The outcome of the comparison is shown in Table 38.

Table 38 - Objective 2 Summary of Estimation Method Comparison Outcomes

Method	Method Description	Average MAPE	Average MBE	Usage Cost
M1-V1	Minimum information using set percentage of asset usage	159%	150%	£10
M1-V2	Utilise baseload calculation with set usage percentage	60%	-33%	£30
M2	Baseline comparison using cluster analysis	56%	-16%	£30
M3	Regression analysis utilising historical DSR event outcomes	39%	-12%	£80
M4	Building energy modelling	51%	-6%	£180

The results of the second objective help improve knowledge about how existing DSR estimation methods compare with the following key findings: (1) M1-V1 has the highest error level of 159% with the lowest cost to run of £10; (2) M3 has the lowest error level of 39% with the second highest cost to run of £80; (3) the M1 variations have low input requirements through usage of half-hourly site electricity usage data, which reduces analysis time and cost, but causes higher error levels; (4) M2 also only requires half-hourly usage data, yet the method will only work on electrical assets that have enough variation within the building's overall usage to be identified by the clustering; (5) while M3 had the lowest error level, its usage is restricted as it requires the site to have previously undertaken DSR and to be able to provide the outcomes of those DSR events; (6) M4 had the highest cost at £180 due to the amount of time required to develop a building energy model for each site being assessed; (7) sensitivity analysis showed that caution needs to be taken on ensuring accuracy of input value selection, as the outcomes for all methods show sensitivity to the values used.

6.1.3 Objective 3 Review - Development of a Profile Based DSR Estimation Method

Chapter 5 addressed the third research objective *'To develop and evaluate a model that uses asset usage profiles to reduce the uncertainty of DSR potential estimation during an aggregator's assessment process'*. The purpose of this objective was to create a new DSR estimation method that utilised existing research on building energy estimation approaches to improve new site assessments for DSR. A review of building energy usage estimation approaches in section 2.4 identified the load profile method as offering the greatest opportunity for a new DSR estimation method due to: (1) its extensive utilisation for understanding electricity usage; (2) no previous use for site level DSR estimation; (3) its capability for use to estimate variable and static electricity usage levels.

The development of the new profile based DSR estimation method was undertaken using a seven-stage process, as outlined in Figure 16. During the development, three methodology questions were identified and addressed with the outcomes providing important guidance for future load profile creations, both for DSR estimation and the overall body of knowledge of DSR profiling. The first question concerned how to manage missing half-hourly usage values, which was answered by testing different gap sizes and found that fixing single missing gaps resulted in the highest error level reduction. The second question looked at which of the three testing and training dataset selection methods should be used. After testing each method, the '*K-fold Cross Validation Random Selection'* option was identified as providing the most balanced outcome while being the least susceptible to influence from uncharacteristic usage data. The third question related to determining what stochastic sample size should be used when creating the profile distributions from the clusters. As the size directly impacted computation time for creating the profiles, the final size of 3000 samples was chosen. After this sample size, only marginal improvements in error levels occurred, and at the expense of significant increases in processing time.

The third objective results validated that using load profiles for DSR estimation provided many benefits based on the following key findings: (1) once the profiles have been created they can be applied using only three inputs: business category (e.g. hotel), asset type (e.g. HVAC chiller), and the assets maximum electricity usage rating (e.g. 200kW); (2) as the generated profiles show usage variation across the day and between weeks, months or days (depending on the selected predictor) will enable the estimation outputs to be tailored for the flexibility of the intended DSR programme; (3) the profiles' median, quartile and confidence interval usage values enable the potential variance of the estimation output to be measured to understand the level of uncertainty; (4) the uncertainty level information allows users to adjust the estimation outputs based on their risk appetite, a feature not available in the existing four DSR estimation methods; (5) comparison of the profile method in a deterministic form (e.g. only using the median level usage values) against the existing DSR estimation methods found that the new method achieves the second lowest MAPE of 46.5%, and second lowest per usage cost of £26.

6.2 Research Implications

The implications of this research start with the specific topic of DSR estimations, before progressing towards the broader topics of DSR aggregators, DSR uptake and general implications of this research for usage of load profiles.

6.2.1 DSR Estimation Implications

The general body of knowledge on DSR is extensive, as seen in section 2.1. Yet as discovered in section 2.3, there is only limited research on the specific sub-topic of DSR estimation methods. The focus of this thesis on DSR estimation and the limited existing research in this domain means that its outcomes offer three direct contributions to the DSR estimation knowledge base.

The first contribution arises from the review of KiWi Power's aggregator assessment process, described in chapter 3. This review identified that during the assessment of a new site, the tasks which focused on understanding a site's DSR potential are treated by the aggregator as having the highest importance for deciding site suitability for DSR. These tasks include the DSR estimation activities performed once the site enters the second stage of the assessment process, which currently requires obtaining detailed information about the site's half-hourly electricity usage records for a year and assets that could be used for DSR. This is an important finding for DSR research as it helps address the knowledge gap on the role that DSR estimation plays during the suitability assessment of a new site.

The second contribution is the outcomes of chapter 4's DSR estimation methods comparison. Based on a review of existing literature in section 2.3, the review and comparison of existing DSR estimations methods as described in chapter 4 have not been previously undertaken. Therefore, the outcomes from this review have three key implications: (1) enabling current or future users of the DSR estimation methods described to understand the level of estimation uncertainty and sensitivity of inputs of their selected approach; (2) provides knowledge about DSR estimation methods that enables users to select and implement a new method with a full understanding of the input requirements and potential output error levels; (3) consolidates knowledge about existing DSR estimation methods into one location that reduces effort and time for future researchers of this topic.

The third contribution is chapter 5's development of the new load profile based DSR estimation method. The new method offers the ability for aggregators to understand as an output the levels of uncertainty at the time of estimation. This could potentially help improve the decision making about suitability of new sites for DSR by providing an understanding for aggregators of the extent of uncertainty at the time of estimation, which can then be used to adjust with greater accuracy the estimates about risks of a site failing to meet a DSR programme's criteria. The creation of this new method has three implications: (1) it expands the range of available DSR estimation methods; (2) it can improve the suitability assessments undertaken by DSR aggregators as discussed in the following section 6.2.2; (3) it provides additional insights into creation and usage of load profiles, as discussed in section 6.2.4.

6.2.2 DSR Aggregator Implications

Existing research on DSR aggregators is very limited as outlined in section 2.2 and focuses on the issues that aggregators face with integrating into the traditional electricity market. No research was found that reviewed how aggregators operate, which is a significant gap considering their role in providing over 80% of DSR capacity. This means that chapter 3's research into how an aggregator (in this case KiWi Power) operates is the first of its kind, and provides valuable initial insights into the DSR suitability assessment process for new sites.

A key insight from the aggregator process review was the continual importance placed by the aggregator on tasks that aim to assess the new site's DSR potential. During the first stage telephone contact, for example, the aggregator's business development team prioritise trying to determine what potential assets could be used for DSR, placing less importance on understanding any impacts from DSR programmes on the site's business processes. The asset assessment during the first stage is at a high level, and relies on the site contact providing verbal answers to questions about the kW size of generators and large electricity using systems (like HVAC chillers). It is not until the second stage of a site review that detailed information about the assets is obtained and analysed using DSR estimation methods to determine a site's DSR potential. As a result, any decision to reject a site during the first contact is solely based on perceived asset potential, and made using limited information that could lead to incorrect assessments of a site's true potential for DSR programmes.

However, the aggregator does not complete desktop DSR estimation until the second stage of an assessment due to the time and effort involved in obtaining a site's half-hourly electricity records, which is required to undertake the estimation using method 1 as described in section 6.1.2). An implication of the new profile DSR estimation method outlined in chapter 5 is that it could be used during the first stage and provide additional analysis support for DSR suitability decision-making. The new method could be used by an aggregator even if the existing ones cannot because the new profile method proposed by this research only requires three information inputs to run: (1) the business category (e.g. hotel), (2) the asset type (e.g. HVAC chiller), and (3) the asset's maximum kW usage rating (e.g. 200kW). All these inputs are currently obtained during the stage 1 assessment process, with the business category being known before the call and the assets' types and ratings being obtained during the initial call.

Therefore, an initial DSR estimation could be undertaken using the new profile method during first contact with a prospective site if profiles have been created for the business category and assets in use at a site. Undertaking this assessment during stage 1 might result in two important implications. It could help overcome the current top three reasons for clients being lost during this first stage of *'no interest in DSR'* at 37%, *'No assets or generators large enough for DSR'* at 19%, and '*Assets or generators deemed unsuitable for DSR'* at 14% as using load profiles would enable the sales person to provide a DSR estimation during the initial call, which could overcome any lack of initial interest by highlighting the potential financial returns or verifying that the site's assets are sufficient to process to stage 2. Alternatively, it could result in losing more sites at this first stage if the initial estimate is lower than expected. Yet this could still be a benefit, as the sites that are eliminated during stage one might represent the 35% that formerly made it through stage 1 but then dropped out during stage 2, suggesting that filtering out the maximum quantity of sites during stage 1 saves effort (and time and cost) for both parties. It could also be a disadvantage if the estimate is lower than the site owner's own expectations, and results in loss of interest at the start. If a site had instead progressed through to stage 2, as it might with the current method, then the additional contact time with the aggregator could help improve interest and the strength of the relationship with the aggregator, potentially ultimately generating appetite for seeking more opportunities for DSR at the site even if stage one estimates for potential are quite low. However, this is more about relationship management and selling skills, than an issue with the capabilities of the new method.

6.2.3 DSR Uptake Implications

The primary motivation for this research is enabling insights capable of contributing towards an increase in the uptake of DSR. The previous two sections address the primary ways that this research contributes towards increasing DSR uptake, through improving the understanding of DSR estimation methods and aggregator suitability assessments for new sites. This next section reviews two additional implications of the research findings on the uptake of DSR.

The first additional DSR uptake implication emerges from the chapter 3 finding that the aggregator's sales team tend to require a minimum turndown potential of between 50kW to 100kW per asset and 200kW for a site to be deemed financially worth progressing for any DSR programme. Interviews with the sales team found that these values were not officially defined or mandatory company policy, and more based on what staff had been told anecdotally or independently determined based on their own experiences. This implies a size barrier affecting DSR uptake for two reasons. Firstly, the current UK DSR programmes as outlined in section 2.1.1 do not obviously provide sufficient financial incentives for aggregators and their sales staff to use sites with a DSR potential below these levels. Secondly, the cost of enablement is too expensive to justify any investment in assets below these minimum turndown levels. Both causes could be overcome by DSR market providers (for example, the UK National Grid) introducing new DSR programmes with higher financial incentives for smaller turndown loads. This approach is unlikely to occur without financial assistance from government targeted at encouraging uptake, as it seems likely that existing programmes would already offer these variants for existing programmes if it made financial sense. More realistically this issue will need to be addressed by reducing enablement costs, if small turndown loads are to become financially viable. Cost reduction could be achieved by improving processes and/or lowering technical solution costs.

The experience of KiWi Power's current processes indicates that there are limited improvements to be made by this aggregator without fundamental changes to how they operate and engage with customers, as the company has already refined their process to reduce costs where possible. By way of example, a fundamental change to reduce costs might include shifting from the current approach of the aggregator's sales teams contacting end users to assess suitability to a new approach, whereby end users directly access a self-serving internet-based system to perform the assessment themselves. End users would then undertake their own site analysis and qualification via estimation tools available on the website, potentially also then ordering and receiving the necessary monitoring and control equipment that the users could themselves organise to have installed. This would significantly reduce the aggregator's costs of acquiring new sites. However, it would require significant upfront costs to develop, be technically challenging to provide a monitoring and control system capable of use with a wide range of assets and minimum end user configuration, could increase the potential for unsuitable sites if end users use incorrect information, and would require end users to proactively seek out DSR opportunities. Any successful outcome from the aggregator from adopting self-service for suitability assessments would also rely on investment by the aggregator in proactive marketing and communications programmes, if end users are to choose the aggregator's web-based service, over other companies.

Another option for fundamental change is to reduce the technical solution costs for enabling smaller sites. Currently KiWi Power's monitoring and control hardware costs approximately £300 per asset, plus a minimum of £500 for a contractor to install. While this cost is expected to decrease as the hardware is refined, it may still be prohibitive in relation to the returns from assets with small turndown potential, and require costly contractors to install. To overcome these costs would either require new hardware that is cheap and easy to install, or the assets come with DSR functionality built in that can be enabled by the end user with minimum involvement.

The second DSR uptake implication for this research results from chapter 3, which found that only 36% of sites that started the aggregator's assessment completed the process and went live with DSR. 72% of the time a decision not to progress was made by the end user, with the primary reason for ending the process being 'loss of interest', which occurred 35% of the time. As this is the first known research that assess an aggregator's operations, the outcomes from chapter 3 provide a new perspective on the reasons why end users do not progress with DSR. While the majority of the reasons emerging from this research align with similar end user surveys by The Energyst (2017) and Olsthoorn et al. (2015), the high 'loss of interest' reason stands out, as The Energyst results for this reason category were a third lower at 23%.

Additional detail on why the end users lost interest could not be obtained from the KiWi Power's available data. However, it is not apparent that this loss of interest is related to operational impact, financial, or asset suitability as these reasons for eliminating a site are captured separately by the aggregator during its assessment process. Addressing the reasons why end users may decide not to participate is important if uptake is to increase. Therefore, this topic offers an opportunity for potential future research work, as outlined in section 6.3 overcoming this lack of interest will be important to increasing the uptake of DSR. Doing so may require multiple paths to address, and potentially include: (1) education on DSR benefits; (2) government regulation to encourage DSR usage; (3) improved financial benefits; and (4) new DSR programmes that are more suitable for end users. The actual paths that need to be taken cannot be verified until additional research on the lack of interest is undertaken to ascertain its underlying causes.

6.2.4 Load Profile Method Implications

The development of the new profile based DSR estimation in chapter 5 was based on previous research undertaken on usage of load profiles for estimation energy usage in buildings. The impact of using load profiles for DSR estimation has been reviewed in the previous sections 6.2.1 to 6.2.3. However, it also has implications for the general non-DSR analysis of load profiles. The first implication arises from the testing and selection of approaches used during development of the profile method, as detailed in section 5.8. At a high level, the load profile literature reviewed in section 5 all use the same general steps for creating profiles. This entails first obtaining multiple usage datasets on the subject area, grouping usage by similarity, and then converting the groups into common profiles. However, the literature about load profile usage for estimating general energy usage in buildings also showed that different methods were used to achieve each step. As a consequence, the development stages of the DSR profile estimation method outlined in this research involved having to also test the available methods to ascertain the optimal approach in adopting load profiles for DSR estimation. The method selections used during this development have implications for future load profile research and creation, as the analysis undertaken and options selected provide additional supporting evidence of the methods adopted in earlier research, which may assist future researchers reviewing options for creating their own load profiles.

The second implication of this research for use of load profiles in estimating energy usage is how uncertainty in the profiles can be determined at the time of usage, by creating multiple percentile-based usage levels for each profile. These levels can then be used to understand the level of potential variation in the estimation outcomes, and resulting likely variations from actual energy usage. This application of percentile-based usage levels for the profiles differs from the existing literature reviewed in section 5, as existing research and application of load profiles only provide one usage level based on an averaged value of the grouped data. A key example of this is the Elexon profile classes, which are used to determine expected usage for domestic and non-domestic sites that lack half-hourly metering in the UK (Elexon, 2013). Each Elexon profile is created using the averaged values of the sample usage groups. This results in no ability to understand variation within the profiles which in turn leads to increased uncertainty of outcomes. Therefore, the ability to understand the estimation outcome variation arising from the application of load profiles, as demonstrated in this thesis, enables uncertainty levels to be known at the time of a site's estimation and factored into decision-making processes.

6.3 Future Research

Throughout this research a number of limitations were encountered that offer the potential for future research opportunities. The first research opportunity was found from chapter 3's review of the aggregator's new site assessment process. One of the main reasons that sites decided not to proceed with DSR is 'loss of interest'. The research was unable to ascertain the reasons why end users lost interest from the available data. Understanding why end users lost interest in DSR during suitability assessments would enable solutions to be developed to help overcome this uptake barrier. However, future research in this area could encounter several impediments to successful completion where commercial or privacy restrictions preclude access to any historical details about the end users, or there's low levels of confidence that end users who are contactable will be able to accurately recall and describe their reasons for not proceeding. These barriers could be overcome if this research is undertaken by the aggregator's sales team on each occasion that an end user indicates a wish not to proceed further due to 'lack of interest'.

The chapter 3 review of an aggregator's new site assessment process was also constrained by only having access to one aggregator, KiWi Power. This means that the outcomes of this review are based on KiWi Power's experience, and while it is assumed that KiWi Power's operations are like other aggregators, this cannot be verified without undertaking the same review using a different aggregator. This offers another future research opportunity for performing the same analysis as undertaken in chapter 3 on different aggregators and then comparing the outcomes to understand how these may differ. This could enable greater understanding of how sites are assessed and provide further insights into why sites are lost during the assessment process. However, this research could be problematic to undertake due to commercial sensitivities, and while another aggregator might be willing to be assessed and happy to review the results internally, it is anticipated that there could be concerns with publishing the results.

A limitation of the research informing this thesis was the availability of suitable data, which meant that the dataset used in chapters 4 and 5 was necessarily limited to HVAC chiller usage in UK hotels. This limitation provides the third and fourth research opportunities by extending the work done in chapters 4 and 5 through usage of new datasets. The comparison of DSR estimation methods in chapter 4 was undertaken using data obtained from two hotels. Future research could re-run this comparison using data from a different business category, assets or both variables to understand how different businesses and assets impact estimation outcomes and error levels. Similar, more hotel HVAC chiller data could be added to chapter 5's profile method to assess how this impacts the resulting profiles. Additionally, creating profiles using data from different assets and businesses would enable a comparison of the outcomes to understand how profiles are impacted by different types of usage data. A recommended additional dataset would be from HVAC chillers in offices, as an example that would allow for a direct comparison with the results from this research based on only changing the business category input. *The end.*

7 References

ADE. (2018). What is demand side response. Retrieved August 7, 2018, from https://www.theade.co.uk/resources/what-is-demand-side-response

Aghabozorgi, S., Seyed Shirkhorshidi, A., & Ying Wah, T. (2015). Time-series clustering - A decade review. *Information Systems*, *53*, 16–38. https://doi.org/10.1016/j.is.2015.04.007

Agora Energiewende. (2013). Load Management as a Way of Covering Peak Demand. Retrieved July 18, 2017, from www.agora-energiewende.de

Aman, S., Frincu, M., Chelmis, C., Noor, M., Simmhan, Y., & Prasanna, V. K. (2016). Prediction models for dynamic demand response: Requirements, challenges, and insights. *2015 IEEE International Conference on Smart Grid Communications, SmartGridComm 2015*, 338–343. https://doi.org/10.1109/SmartGridComm.2015.7436323

Andrews, D. (2007). The Potential Contribution of Emergency Diesel Standy Generators. In *Renewable Energy and the Grid: The Challenge of Variability* (p. 219). Enfield: Earthscan.

ASHRAE. (2014). *Guideline 14-2014 -- Measurement Of Energy, Demand, And Water Savings*. ASHRAE.

ASHRAE. (2015). *M&V Guidelines: Measurement and Verification for Performance-Based Contracts - Version 4.0*. Retrieved from http://www1.eere.energy.gov/

Barton, J., Huang, S., Infield, D., Leach, M., Ogunkunle, D., Torriti, J., & Thomson, M. (2013). The evolution of electricity demand and the role for demand side participation, in buildings and transport. *Energy Policy*, *52*, 85–102. https://doi.org/10.1016/j.enpol.2012.08.040

Berger, L. T., & Iniewski, K. (2012). *Smart Grid Applications, Communications, and Security*. Hoboken: Wiley.

Bergmeir, C., Hyndman, R. J., & Koo, B. (2015). *A Note on the Validity of Cross-Validation for Evaluating Time Series Prediction*. *Monash University, Department of Econometrics and Business Statistics*.

Better Building Partnership. (2011). *Better Metering Toolkit*. Retrieved from http://www.betterbuildingspartnership.co.uk/download/bbp-better-metering-toolkit.pdf

Blumsack, S., & Fernandez, A. (2012). Ready or not, here comes the smart grid! *Energy*, *37*(1), 61–68. https://doi.org/10.1016/j.energy.2011.07.054

Borgstein, E. H., Lamberts, R., & Hensen, J. L. M. (2016). Evaluating energy performance in non-domestic buildings: A review. *Energy and Buildings*, *128*, 734–755. https://doi.org/10.1016/j.enbuild.2016.07.018

Bryman, A. (2016). *Social Research Methods* (5th ed.). Oxford: Oxford University Press.

Building Energy Software Tools. (2017). Building Energy Software Tools. Retrieved August 7, 2017, from www.buildingenergysoftwaretools.com

Burger, S., Chaves-Ávila, J. P., Batlle, C., & Pérez-Arriaga, I. J. (2017). A review of the value of aggregators in electricity systems. *Renewable and Sustainable Energy Reviews*, *77*(February), 395–405. https://doi.org/10.1016/j.rser.2017.04.014

Burman, E., Hong, S.-M., Paterson, G., Kimpian, J., & Mumovic, D. (2014). A comparative study of benchmarking approaches for non-domestic buildings: Part 2 – Bottom-up approach. *International Journal of Sustainable Built Environment*, *3*(2), 247–261. https://doi.org/https://doi-org.idpproxy.reading.ac.uk/10.1016/j.ijsbe.2014.12.001

California Energy Commission. (2006). California Commercial End-Use Survey (CEUS). *Analysis*, (March).

Cappers, P., MacDonald, J., Goldman, C., & Ma, O. (2013). An assessment of market and policy barriers for demand response providing ancillary services in U.S. electricity markets. *Energy Policy*, *62*, 1031–1039. https://doi.org/10.1016/j.enpol.2013.08.003

Carrie, K., Gupta, A., Shrimali, G., & Albert, A. (2013). Is disaggregation the holy grail of energy efficiency?

The case of electricity. *Energy Policy*, *52*, 213–234. https://doi.org/10.1016/j.enpol.2012.08.062

Chen, Y., Xu, P., Chu, Y., Li, W., Wu, Y., Ni, L., … Wang, K. (2017). Short-term electrical load forecasting using the Support Vector Regression (SVR) model to calculate the demand response baseline for office buildings. *Applied Energy*, *195*, 659–670. https://doi.org/10.1016/j.apenergy.2017.03.034

Chicco, G. (2012). Overview and performance assessment of the clustering methods for electrical load pattern grouping. *Energy*, *42*(1), 68–80. https://doi.org/10.1016/j.energy.2011.12.031

Christoph, B., & José, B. (2012). On the use of cross-validation for time series predictor evaluation. *Information Sciences*, *191*, 192–213. https://doi.org/10.1016/j.ins.2011.12.028

CIBSE. (2006a). *CIBSE TM22 Energy Assessment and Reporting Methodology*.

CIBSE. (2006b). *TM41: Degree Days: Theory & Application*. London: CIBSE Publications.

CIBSE. (2012a). *CIBSE Guide F*. London: CIBSE Publications.

CIBSE. (2012b). *CIBSE Guide F - Energy efficiency in buildings* (3rd ed.). Lavenham,: The Lavenham Press Ltd.

CIBSE. (2015). *AM11 Building Performance Modelling*.

Clarke Energy. (2016). Island Mode Operation Captive Power Plant. Retrieved June 8, 2016, from https://www.clarke-energy.com/gas-engines/island-mode-operation/

Coakley, D., Raftery, P., & Keane, M. (2014). A review of methods to match building energy simulation models to measured data. *Renewable and Sustainable Energy Reviews*, *37*(Supplement C), 123–141. https://doi.org/https://doi-org.idpproxy.reading.ac.uk/10.1016/j.rser.2014.05.007

Commons Select Committees. (2015). *Energy and Climate Change Committee - Eighth Report: Implementation of Electricity Market Reform*. Retrieved from https://publications.parliament.uk/pa/cm201415/cmselect/cmenergy/664/66406.htm

Curtis, M., Torriti, J., & Smith, S. T. (2018). Demand Side Flexibility and Responsiveness: Moving Demand in Time Through Technology. In A. Hui, R. Day, & G. Walker (Eds.), *Demanding Energy: Space, Time and Change* (pp. 283–312). Cham: Springer International Publishing. https://doi.org/10.1007/978-3-319-61991-0_13

De Gooijer, J. G., & Hyndman, R. J. (2006). 25 Years of Time Series Forecasting. *International Journal of Forecasting*, *22*(3), 443–473. https://doi.org/10.1016/j.ijforecast.2006.01.001

DECC. (2013). Electricity Market Reform (EMR). Retrieved December 15, 2014, from https://www.gov.uk/government/policies/maintaining-uk-energy-security--2/supporting-pages/electricity-market-reform

DesignBuilder. (2017a). DesignBuilder - Edit / Translate Hourly Weather Data. Retrieved February 3, 2017, from http://www.designbuilder.co.uk/helpv5.0/Content/_Edit_hourly_weather_data.htm

DesignBuilder. (2017b). DesignBuilder V5.0.2. Retrieved from https://designbuilder.co.uk/software/product-overview

Dronkelaar, C., Dowson, M., Burman, E., Spataru, C., & Mumovic, D. (2016). A Review of the Energy Performance Gap and Its Underlying Causes in Non-Domestic Buildings. *Frontiers in Mechanical Engineering*, *1*, 17. https://doi.org/10.3389/fmech.2015.00017

DRRC. (2017). Demand Response Research Center. Retrieved August 5, 2016, from https://drrc.lbl.gov/about-us

Dudley, J. H. (2010). Comparison of Demand Response Performance with an EnergyPlus Model in a Low Energy Campus Building. Retrieved from http://www.escholarship.org/uc/item/9zw4m12g

ECMWF. (2017). European Centre for Medium-Range Weather Forecasts - Meteorological archive. Retrieved May 2, 2017, from https://www.ecmwf.int/en/forecasts/accessing-forecasts/order-historical-datasets

EIA. (2017). *Trends in Lighting in Commercial Buildings*. Retrieved from https://www.eia.gov/consumption/commercial/reports/2012/lighting/

Element Energy. (2012). Demand side response in the non-domestic sector. Retrieved January 16, 2014, from https://www.ofgem.gov.uk/ofgem-publications/57014/demand-side-response-non-domestic-sector.pdf

Elexon. (2013). *Load Profiles and their use in Electricity Settlement*. Retrieved from http://www.elexon.co.uk/wp-content/uploads/2013/11/load_profiles_v2.0_cgi.pdf

Elexon. (2014). Metering. Retrieved December 5, 2014, from http://www.elexon.co.uk/reference/technical-operations/metering/

European Commission. (2007). 2020 Climate & Energy Package. Retrieved February 22, 2017, from https://ec.europa.eu/clima/policies/strategies/2020_en

Furnham, A. (1986). Response bias, social desirability and dissimulation. *Personality and Individual Differences*, *7*(3), 385–400. https://doi.org/10.1016/0191-8869(86)90014-0

Gao, X., & Malkawi, A. (2014). A new methodology for building energy performance benchmarking: An approach based on intelligent clustering algorithm. *Energy & Buildings*, *84*, 607–616. https://doi.org/10.1016/j.enbuild.2014.08.030

Gkatzikis, L., Koutsopoulos, I., & Salonidis, T. (2013). The role of aggregators in smart grid demand response markets. *IEEE Journal on Selected Areas in Communications*, *31*(7), 1247–1257. https://doi.org/10.1109/JSAC.2013.130708

Grein, A., & Pehnt, M. (2011a). Load management for refrigeration systems: Potentials and barriers. *Energy Policy*, *39*(9), 5598–5608. https://doi.org/10.1016/j.enpol.2011.04.040

Grein, A., & Pehnt, M. (2011b). Load management for refrigeration systems: Potentials and barriers. *Energy Policy*, *39*(9), 5598–5608. https://doi.org/10.1016/J.ENPOL.2011.04.040

Grid, N. (2017). What is Power Responsive? Retrieved June 18, 2017, from http://powerresponsive.com/faqs/#demand

Grünewald, P., & Torriti, J. (2013). Demand response from the non-domestic sector: Early UK experiences and future opportunities. *Energy Policy*, *61*, 423–429. https://doi.org/10.1016/j.enpol.2013.06.051

Hadjsaid, N., & Sabonnadiere, J.-C. (2013). *Smart Grids*. London: iSTE Ltd.

Harpe, S. E. (2015). How to analyze Likert and other rating scale data. *Currents in Pharmacy Teaching and Learning*, *7*(6), 836–850. https://doi.org/10.1016/j.cptl.2015.08.001

Hartley, J. (2014). Some thoughts on Likert-type scales. *International Journal of Clinical and Health Psychology*, *14*(1), 83–86. https://doi.org/10.1016/S1697-2600(14)70040-7

Hasson, D., & Arnetz, B. B. (2005). Validation and Findings Comparing VAS vs. Likert Scales for Psychosocial Measurements. *International Electronic Journal of Health Education*, *8*, 178–192. https://doi.org/10.1186/1471-2458-5-78

Hay, R., & Macwhinnie, N. (2015). *Making the electricity system more flexible and delivering the benefits for consumers: Position paper*. London. Retrieved from https://www.ofgem.gov.uk/sites/default/files/docs/2015/09/flexibility_position_paper_final_0.pdf

Hoare, J. (2015). *Export Limiting in Practice*. Retrieved from http://www.bre.co.uk/filelibrary/events/BRE Events/Jonny Williams/10-James-Hoare--2-.pdf

Hsu, D. (2015). Comparison of integrated clustering methods for accurate and stable prediction of building energy consumption data. *Applied Energy*, *160*, 153–163. https://doi.org/10.1016/j.apenergy.2015.08.126

Imran. (2015). What is “Within cluster sum of squares by cluster” in K-means. Retrieved March 2, 2017, from https://discuss.analyticsvidhya.com/t/what-is-within-cluster-sum-of-squares-by-cluster-in-k-means/2706/2

IRGC. (2015). *Demand-side flexibility for energy transitions*.

ISO. (2017). *ISO 52016-1:2017 Energy performance of buildings -- Energy needs for heating and cooling, internal temperatures and sensible and latent heat loads*.

Jack, M. W., Suomalainen, K., Dew, J. J. W., & Eyers, D. (2018). A minimal simulation of the electricity demand of a domestic hot water cylinder for smart control. *Applied Energy*, *211*(November 2017), 104–112. https://doi.org/10.1016/j.apenergy.2017.11.044

Jones, P. (2012). *Problems with Sub-Metering Energy Consumption In UK*. *CIBSE Technical Symposium*.

Kalogirou, S. A. (2015). Artificial neural networks in energy. *International Journal of Low Carbon Technologies*, (October), 201–216.

KDnuggets. (2017). New Leader, Trends, and Surprises in Analytics, Data Science, Machine Learning Software Poll. Retrieved July 3, 2017, from https://www.kdnuggets.com/2017/05/poll-analytics-data-science-machine-learning-software-leaders.html

Ketchen, D., & Shook, C. (1996). The application of cluster analysis in strategic management research: An analysis and critique. *Strategic Management Journal*, *17*(6), 441–458. https://doi.org/10.1002/(SICI)1097-0266(199606)17:6<441::AID-SMJ819>3.0.CO;2-G

Kim, Y.-J., Yoon, S.-H., & Park, C.-S. (2013). Stochastic comparison between simplified energy calculation and dynamic simulation. *Energy and Buildings*, *64*(Supplement C), 332–342. https://doi.org/https://doi-org.idpproxy.reading.ac.uk/10.1016/j.enbuild.2013.05.026

Laerd. (2017). Measures of Central Tendency. Retrieved April 2, 2017, from https://statistics.laerd.com/statistical-guides/measures-central-tendency-mean-mode-median.php

Larsen, E. M., Pinson, P., Leimgruber, F., & Judex, F. (2015). From demand response evaluation to forecasting - Methods and results from the EcoGrid EU experiment. *Submitted to IEEE Transactions on Power Systems*, *10*, 75–83. https://doi.org/10.1016/j.segan.2017.03.001

Lau, E., & Livina, V. (2015). *Carbon savings of demand side response of a UK energy aggregator*. *NPL REPORT MS 19*.

Li, Z., Han, Y., & Xu, P. (2014). Methods for benchmarking building energy consumption against its past or intended performance: An overview. *Applied Energy*, *124*, 325–334. https://doi.org/https://doi.org/10.1016/j.apenergy.2014.03.020

Liddiard, R. (2014). Room-scale profiles of space use and electricity consumption in non-domestic buildings. *Building Research & Information*, *42*(1), 72–94. https://doi.org/10.1080/09613218.2013.817112

Lippert, J. (2009). You Can't Manage Energy Use That You Don't Measure. Retrieved February 21, 2017, from https://energy.gov/energysaver/articles/you-cant-manage-energy-use-you-dont-measure

Lockheed Martin Aspen. (2006). Demand Response Enabling Technologies For Small-Medium Businesses. Retrieved May 29, 2016, from http://sites.energetics.com/MADRI/pdfs/LMADRT_060506.pdf

Mathieu, J. L., Gadgil, A. J., Callaway, D. S., Price, P. N., & Kiliccote, S. (2010). Characterizing the Response of Commercial and Industrial Facilities to Dynamic Pricing Signals From the Utility. *ASME 2010 4th International Conference on Energy Sustainability, Volume 1*, 1019–1028. https://doi.org/10.1115/ES2010-90266

May, R., Dandy, G., & Maier, H. (2011). Review of Input Variable Selection Methods for Artificial Neural Networks. *Artificial Neural Networks - Methodological Advances and Biomedical Applications*, (August 2016), 362. https://doi.org/10.5772/644

Meijering, E. (2002). A Chronology of Iinterpolation: From Ancient Astronomy to Modern Signal and Image Processing. *Proceedings of the IEEE*, *90*(3), 319–342. https://doi.org/10.1109/5.993400

Menberg, K., Heo, Y., & Choudhary, R. (2016). Sensitivity analysis methods for building energy models: Comparing computational costs and extractable information. *Energy and Buildings*, *133*, 433–445. https://doi.org/10.1016/j.enbuild.2016.10.005

Menezes, A. C., Cripps, A., Bouchlaghem, D., & Buswell, R. (2012). Predicted vs. actual energy performance of non-domestic buildings: Using post-occupancy evaluation data to reduce the performance gap. *Applied Energy*, *97*, 355–364. https://doi.org/10.1016/j.apenergy.2011.11.075

Menke, R., Abraham, E., Parpas, P., & Stoianov, I. (2016). Demonstrating demand response from water distribution system through pump scheduling. *Applied Energy*, *170*, 377–387. https://doi.org/10.1016/J.APENERGY.2016.02.136

Merry, J. (2017). Inferior submetering is destroying performance. Retrieved April 8, 2017, from https://www.cibsejournal.com/opinion/inferior-submetering-means-substandard-performance/

Mudie, S., Essah, E. A., Grandison, A., & Felgate, R. (2014). Electricity use in the commercial kitchen. *International Journal of Low-Carbon Technologies*, *11*(1), 66–74. https://doi.org/10.1093/ijlct/ctt068

Nadel, S., & Geller, H. (1996). Utility DSM: What have we learned? Where are we going? *Energy Policy*, *24*(4), 289–302. Retrieved from http://www.sciencedirect.com/science/article/pii/0301421595001379

National Grid. (2015). STOR Fuel Type Analysis Summary (By Capacity) - Season 8_5. Retrieved May 30, 2017, from http://www2.nationalgrid.com/WorkArea/DownloadAsset.aspx?id=37710

National Grid. (2016a). *A short guide to how your business can profit from Demand Side Response*. Retrieved from http://powerresponsive.com/wp-content/uploads/pdf/Power Responsive Guide - v8.pdf

National Grid. (2016b). Demand-Side Response Market Snapshot.

National Grid. (2016c). Short Term Operating Reserve Information. Retrieved January 30, 2017, from http://www2.nationalgrid.com/UK/Services/Balancing-services/Reserve-services/Short-Term-Operating-Reserve/Short-Term-Operating-Reserve-Information/

National Grid. (2016d). Triad. Retrieved November 14, 2017, from www2.nationalgrid.com/WorkArea/DownloadAsset.aspx?id=45740

National Grid. (2017a). *Firm Frequency Response FAQ*. Retrieved from http://www2.nationalgrid.com/UK/Services/Balancing-services/Frequency-response/Firm-Frequency-Response/FFR-Bridging/

National Grid. (2017b). Frequency Control by Demand Management. Retrieved September 9, 2017, from http://www2.nationalgrid.com/uk/services/balancing-services/frequency-response/frequency-control-by-demand-management/

National Grid. (2017c). GB Electricity Capacity Market Notices. Retrieved October 1, 2017, from https://gbcmn.nationalgrid.co.uk/

National Grid. (2017d). Keeping the Electricity Transmission System in balance. Retrieved June 5, 2016, from http://www2.nationalgrid.com/uk/services/balancing-services/

National Grid. (2017e). *Non-BM Balancing Services Volumes and Expenditure*.

National Grid. (2017f). Short Term Operating Reserve. Retrieved January 21, 2017, from http://www2.nationalgrid.com/uk/services/balancing-services/reserve-services/short-term-operating-reserve/

National Grid. (2017g). Triad Data. Retrieved November 14, 2017, from http://www2.nationalgrid.com/UK/Industry-information/System-charges/Electricity-transmission/Transmission-Network-Use-of-System-Charges/Transmission-Charges-Triad-Data/

Nau, R. (2017). Three types of forecasts: estimation, validation, and the future. Retrieved October 9, 2017, from https://people.duke.edu/~rnau/three.htm

Neto, A. H., & Fiorelli, F. A. S. (2008). Comparison between detailed model simulation and artificial neural network for forecasting building energy consumption. *Energy and Buildings*, *40*(12), 2169–2176. https://doi.org/10.1016/j.enbuild.2008.06.013

Nichols, A. L., & Maner, J. K. (2008). The Good-Subject Effect: Investigating Participant Demand Characteristics. *The Journal of General Psychology*, *135*(2), 151–166.

https://doi.org/10.3200/GENP.135.2.151-166

NIST/SEMATECH. (2012a). Anderson-Darling Test. Retrieved April 5, 2017, from http://www.itl.nist.gov/div898/handbook/eda/section3/eda35e.htm

NIST/SEMATECH. (2012b). Measures of Skewness and Kurtosis. Retrieved April 5, 2017, from http://www.itl.nist.gov/div898/handbook/eda/section3/eda35b.htm

Nolan, S., & O'Malley, M. (2015). Challenges and barriers to demand response deployment and evaluation. *Applied Energy*, *152*, 1–10. https://doi.org/10.1016/j.apenergy.2015.04.083

Ofgem. (2017a). *Annual Report on the Operation of the Capacity Market*.

Ofgem. (2017b). *Upgrading Our Energy System - Smart Systems and Flexibility Plan*.

Olsthoorn, M., Schleich, J., & Klobasa, M. (2015). Barriers to electricity load shift in companies: A survey-based exploration of the end-user perspective. *Energy Policy*, *76*, 32–42. https://doi.org/10.1016/j.enpol.2014.11.015

Owen, G., Ward, J., & Pooley, M. (2012). What Demand Side Services Could Customers Offer? *Paper 3, Sustainabillity First*.

Page, J., Kiliccote, S., Dudley, J. H., Piette, M. A., Koch, E., & Lipkin, P. (2011). Automated Demand Response Technology Demonstration for Small and Medium Commercial Buildings. Retrieved February 6, 2014, from http://drrc.lbl.gov/publications/automated-demand-response-technology-demonstration-project-small-and-medium-commercial-

Palensky, P., & Dietrich, D. (2011). Demand side management: Demand response, intelligent energy systems, and smart loads. *IEEE Transactions on Industrial Informatics*, *7*(3), 381–388. https://doi.org/10.1109/TII.2011.2158841

Panapakidis, I. P., Papadopoulos, T. A., Christoforidis, G. C., & Papagiannis, G. K. (2014). Pattern recognition algorithms for electricity load curve analysis of buildings. *Energy and Buildings*, *73*, 137–145. https://doi.org/10.1016/j.enbuild.2014.01.002

Piette, M. A., Mathieu, J. L., Price, P. N., & Kiliccote, S. (2011). Quantifying changes in building electricity use, with application to demand response. *IEEE Transactions on Smart Grid*, *2*(3), 507–518. https://doi.org/10.1109/TSG.2011.2145010

Pistikopoulos, E., Galindo, A., & Dua, V. (2007). *Multi-Parametric Programming: Theory, Algorithms and Applications, Volume 1*. Weinheim: Wiley-VCH.

Postnote 452. (2014). Electricity Demand-Side Response. Retrieved February 6, 2014, from http://www.parliament.uk/business/publications/research/briefing-papers/POST-PN-452/electricity-demandside-response

Proffitt, E. (2016). Profiting from Demand Side Response. Retrieved May 30, 2016, from http://www.powerresponsive.com/media/1140/ng_meuc-dsr-book.pdf

Qureshi, W. a., Nair, N.-K. C., & Farid, M. M. (2011). Impact of energy storage in buildings on electricity demand side management. *Energy Conversion and Management*, *52*(5), 2110–2120. https://doi.org/10.1016/j.enconman.2010.12.008

Räsänen, T., Voukantsis, D., Niska, H., Karatzas, K., & 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. https://doi.org/10.1016/j.apenergy.2010.05.015

Raslan, R., & Davies, M. (2010). Results variability in accredited building energy performance compliance demonstration software in the UK: an inter-model comparative study. *Journal of Building Performance Simulation*, *3*(1), 63–85. https://doi.org/10.1080/19401490903477386

Research Gate. (2017). Is there an ideal ratio between a training set and validation set? Which trade-off would you suggest? Retrieved October 9, 2017, from https://www.researchgate.net/post/Is_there_an_ideal_ratio_between_a_training_set_and_validati on_set_Which_trade-off_would_you_suggest

Rodriguez, A., Smith, S. T., Kiff, A., & Potter, B. (2016). Small power load disaggregation in office buildings based on electrical signature classification. *IEEE International Energy Conference and Exhibition, EnergyCon*, 1–6. https://doi.org/http://dx.doi.org/10.1109/ENERGYCON.2016.7513984

Royset, J. O. (2013). On sample size control in sample average approximations for solving smooth stochastic programs. *Computational Optimization and Applications*, *55*(2), 265–309. https://doi.org/10.1007/s10589-012-9528-1

Saker, D., Coker, P., Vahdati, M., Millward, S., & Carey, C. (2013). Unlocking the demand response potential from domestic hot water tanks. In *4th Annual TSBE EngD Conference Proceedings*.

Salesforce. (2016). Salesforce. Retrieved from www.salesforce.com

Saltelli, A., Chan, K., & Scott, E. (2008). *Sensitivity Analysis*. Chichester: Wiley.

Sayad, S. (2017). K-Means Clustering. Retrieved March 2, 2017, from http://www.saedsayad.com/clustering_kmeans.htm

Schmidt, M., & Bansal, S. (2017). Energy Disaggregation Methods for Commercial Buildings Using Smart Meter and Operational Data. In *AAAI-17* (pp. 233–239).

Scikit-learn. (2016). K-Means clustering function. Retrieved from http://scikit-learn.org/stable/modules/generated/sklearn.cluster.KMeans.html

SDG&E. (2016). *San Diego Gas & Electric - Capacity Bidding Program*. Retrieved from http://regarchive.sdge.com/tm2/pdf/ELEC_ELEC-SCHEDS_CBP.pdf

SEDC. (2015). *Mapping Demand Response in Europe Today*. *Smart Energy Demand Coalition*. Brussels.

SEDC. (2016). *Explicit and Implicit Demand-Side Flexibility Complementary Approaches for an Efficient Energy System*.

SEDC. (2017). *Explicit Demand Response in Europe Mapping the Markets 2017*.

Shenton Group. (2013). *The G59 Process A Simple Overview*. Retrieved from http://www.shentongroup.co.uk/wp-content/uploads/2013/07/The-G59-Process-A-Simple-Overview-V4.pdf

Siano, P. (2014). Demand response and smart grids - A survey. *Renewable and Sustainable Energy Reviews*, *30*, 461–478. https://doi.org/10.1016/j.rser.2013.10.022

Sluiter, R. (2009). *Interpolation methods for climate data: literature review.*

Stack Overflow. (2017). Is there a rule-of-thumb for how to divide a dataset into training and validation sets? Retrieved October 9, 2017, from https://stackoverflow.com/questions/13610074/is-there-a-rule-of-thumb-for-how-to-divide-a-dataset-into-training-and-validatio

Stern, J. (2014). The British utility regulation model: Its recent history and future prospects. *Utilities Policy*, *31*, 162–172. https://doi.org/10.1016/j.jup.2014.09.005

Strbac, G. (2008). Demand side management: Benefits and challenges. *Energy Policy*, *36*(12), 4419–4426. https://doi.org/10.1016/j.enpol.2008.09.030

The Energyst. (2016). *Demand Side Response: Turning Interia into Inertia*. Retrieved from http://theenergyst.com/

The Energyst. (2017). *Demand-side response: Shifting the balance of power: 2017 Report*. Retrieved from https://theenergyst.com/demand-side-response-battery-storage-report/

Tong, Z., & Zhang, K. M. (2015). The near-source impacts of diesel backup generators in urban environments. *Atmospheric Environment*, *109*, 262–271. https://doi.org/https://doi.org/10.1016/j.atmosenv.2015.03.020

Torriti, J., & Grunewald, P. (2014). Demand Side Response : Patterns in Europe and Future Policy Perspectives under Capacity Mechanisms. *Economics of Energy & Environmental Policy*, *3*(1), 69–87.

U.S. Department of Energy. (2017). EnergyPlus. Retrieved April 3, 2017, from https://energyplus.net/

UK Met Office. (2017a). MIDAS Data Guide. Retrieved February 15, 2017, from https://badc.nerc.ac.uk/data/ukmo-midas/ukmo_guide.html

UK Met Office. (2017b). Seasons. Retrieved November 11, 2017, from https://www.metoffice.gov.uk/learning/seasons

UK Met Office. (2017c). UK and regional temp series. Retrieved from https://www.metoffice.gov.uk/climate/uk/summaries/datasets

UKPN. (2017). Demand Side Response. Retrieved June 18, 2017, from http://innovation.ukpowernetworks.co.uk/innovation/en/research-area/demand-side-response/

Uusitalo, L., Lehikoinen, A., Helle, I., & Myrberg, K. (2015). An overview of methods to evaluate uncertainty of deterministic models in decision support. *Environmental Modelling and Software*, *63*, 24–31. https://doi.org/10.1016/j.envsoft.2014.09.017

Van Dievel, P., De Vos, K., Belmans, R., Van Dievel, P., De Vos, K., & Belmans, R. (2014). Demand response in electricity distribution grids: Regulatory framework and barriers. In *11th International Conference on the European Energy Market (EEM14)* (pp. 1–5). https://doi.org/10.1109/EEM.2014.6861286

Van Wijk, J. J., Van Selow, E. R., Wijk, J. J. Van, & Selow, E. R. Van. (1999). Cluster and calendar based visualization of time series data. *Proceedings 1999 IEEE Symposium on Information Visualization (InfoVis'99)*, 4. https://doi.org/10.1109/INFVIS.1999.801851

Varadarajan, M., & Swarup, K. S. (2007). Solving multi-objective optimal power flow using differential evolution. *Generation, Transmission & Distribution, IET*, *1*(2), 324. https://doi.org/10.1049/iet-gtd

Vardakas, J. S., Zorba, N., & Verikoukis, C. V. (2015). A Survey on Demand Response Programs in Smart Grids : Pricing Methods and Optimization Algorithms. *IEEE COMMUNICATION SURVEYS & TUTORIALS, 17*(1), 152–178.

Wang, Y., & Li, L. (2013). Time-of-use based electricity demand response for sustainable manufacturing systems. *Energy*, *63*, 233–244. https://doi.org/10.1016/j.energy.2013.10.011

Warren, P. (2014). A review of demand-side management policy in the UK. *Renewable and Sustainable Energy Reviews*, *29*, 941–951. https://doi.org/10.1016/j.rser.2013.09.009

Widén, J., Lundh, M., Vassileva, I., Dahlquist, E., Ellegård, K., & 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. https://doi.org/10.1016/j.enbuild.2009.02.013

Willits, F., Theodori, G. L., & Luloff, A. E. (2016). Another look at likert scales. *Rural Social Sciences*, *31*(August 2015), 126–139.

WPD. (2017). Project ENTIRE. Retrieved June 18, 2017, from https://www.westernpower.co.uk/Innovation/Projects/Current-Projects/Project-ENTIRE.aspx

Yao, R., & Steemers, K. (2005). A method of formulating energy load profile for domestic buildings in the UK. *Energy and Buildings*, *37*(6), 663–671. https://doi.org/https://doi-org.idpproxy.reading.ac.uk/10.1016/j.enbuild.2004.09.007

Yin, R., & Black, D. (2015). *Improvement of Demand Response Quick Assessment Tool (Drqat) And Tool Validation Case Studies*.

Zhao, H. X., & Magoulès, F. (2012). *A review on the prediction of building energy consumption*. *Renewable and Sustainable Energy Reviews* (Vol. 16). Elsevier Ltd. https://doi.org/10.1016/j.rser.2012.02.049

Zhou, K. Le, Yang, S. L., & Shen, C. (2013). A review of electric load classification in smart grid environment. *Renewable and Sustainable Energy Reviews*, *24*, 103–110. https://doi.org/10.1016/j.rser.2013.03.023

8 Appendix A

Appendix A contains the interview guide used in chapter 3 for objective 1.

- **Section 0 – Information Sheet and Consent Form for interviewee**
  0. Initial Information provided to interviewee before starting interview with the interview subject number, role type, date and time recorded. Italic usage in this guide indicates the parts that the interviewer will speak to the interviewee.
     0.1. Hi x, thank you for taking the time today to participate in my research project.
     0.2. Before we start could you please read through this information sheet and also please read and sign this Consent Form, let me know if you have any questions about these.
     0.3. I have structured this interview to try and understand the steps and decisions made during each of the major stages of the KiWi Power new client workflow. The main purpose of which, as noted in the information sheet, is to determine where I can help improve the analysis tools to make the process better by providing more accurate results on the potential of new clients. So, I will be trying to understand what steps are taken during each stage and the decisions you make when deciding if a client should be pursued or not. I have split the stages into four areas – Lead, Opportunity, Project, and Live. Some of the stages will be more relevant than others and we can skip questions that are not applicable to your area of the business.

- **Section 1 – Initial Client Contact Stage (Stage 1)**
  - Description:
    The 'Initial Client Contact' stage covers the initial contact between the aggregator and a potential client, this is typically a 'cold call' situation with the aggregator calling the client. In a small number of occasions, the clients may contact the aggregator first to discuss DR. If a client passes the initial selection criteria (as gauged via a phone call) then they are then moved to the next stage.

  1. Initial Client Contact Stage Interview Questions:
     1.0. I will start with the 'Initial Client Contact' stage first, covering from first contact with a client until a decision is made to either turn them into an opportunity or close the lead.

     1.1. (Note – provide interviewee with sheet 'Question 1.1 Additional Information – Lead Stage Steps') *Could you please have a look at the steps outlined on this sheet and see if you agree/disagree with any of the currently listed steps – please made any changes on the sheet. Also, can you please indicate how important each task is when deciding if the site is suitable or not for demand response?*

     1.2. Of the discussed steps, which ones help you decide if the site has sufficient DR potential to progress

     1.3. What is the minimum DR potential needed to progress a site and how does this vary based on programme/assets available?

1.4. (Note – provide interviewee with sheet ‘Question 1.4 Additional Information – Initial Client Contact Closure Reasons’) *I’ve been looking through Salesforce to try and see if I can determine the common reasons that leads are closed and have created the following list.*

1.4.1. Do you agree/disagree with this list, am I missing any that you think should be included?

1.4.2. *How would you rank them in likelihood reason of causing a lead to be closed?*

- **Section 2 – Site Assessment Stage (Stage 2)**
  - Description:

    The ‘Site Assessment’ or ‘Qualification’ Stage covers the tasks undertaken to verify the potential of the client’s site, determine financials, perform a site visit (if needed), undertake contract negotiations, and create a project plan. Once the contract is signed it will then be handed over to the technical team.

2. Site Assessment Stage Interview questions

2.0. *I will now move onto the Site Assessment or Qualification stage, covering from when an initial contact is made into an opportunity until it has been handed over to delivery as a project to be delivered.*

2.1. (Note – provide interviewee with sheet ‘Question 2.1 Additional Information – Site Assessment Stage Steps) *Could you please have a look at the steps outlined on this sheet and see if you agree/disagree with any of the currently listed steps – please made any changes onto the sheet, also, can you please indicate how important each task is when deciding if the site is suitable or not for demand response?*

2.2. Of the discussed steps, which ones help you decide if the site has sufficient DR potential to progress?

2.3. What kinds of issues are normally found at this stage when trying to determine the DR potential of the site?

2.4. (Note – provide interviewee with sheet ‘Question 2.4 Additional Information – Site Assessment Closure Reasons’) *I have also looked through Salesforce to try and see if I can determine the common reasons that opportunities are closed and have created the following list.*

2.4.1. Do you agree/disagree with this list, am I missing any that you think should be included?

2.4.2. How would you rank them in likelihood reason of causing a lead to be closed?

- **Section 3 – Site Enablement and Go Live Stage (Stage 3)**
  - Description:

    The ‘Site Enablement and Go Live’ or ‘Delivery’ Stage covers the tasks undertaken to enable the site to go live with DSR and normally involves the actual installation and testing of control and monitoring equipment. Once completed it then goes ’Live’ and handed over to operations.

3. Site Enablement and Go Live Stage Interview questions

3.0. *I would now like to discuss your involvement in the delivery process after its been signed and handed over from the sales team.*

3.1. (Note – provide interviewee with sheet 'Question 3.1 Additional Information – Site Enablement and Go Live Stage Steps') *Could you please have a look at the steps outlined on this sheet and see if you agree/disagree with any of the currently listed steps – please made any changes onto the sheet. Also, can you please indicate how important each task is when deciding if the site is suitable or not for demand response?*

3.2. *In relation to the DR potential of the site, how often does this turn out to be different to the predicted amount at this stage?*

3.3. *Do changes in the DR potential impact the site going live?*

3.4. *Do you have any examples of sites where the DR potential has changed significantly compared to what was predicted during installation?*

9 Appendix B

Appendix B contains the additional interview information sheets given to interviewees in chapter 3 for objective 1.

Question 1.1 - Additional Information – Initial Client Contact Stage Steps

Please review the following stage steps and:

1. Add any missing steps or adjust any of the currently listed steps.
2. Based on your experience please indicate how important the task is when deciding if the site is suitable for Demand Response. With 1 meaning the task/outcome is not very important and 5 meaning this task/outcome is very important for deciding the suitability of the site.

Task	Initial Client Contact Stage Tasks	How important is this task in determining the suitability of a site for Demand Response?					
		Unimportant				Very Important	N/A
		1	2	3	4	5	
1	Prequalify potential site/client before making contact						
2	Getting in contact with the right person in the company						
3	Explaining Demand Response and checking to see if interested						
4	Checking to see what turndown assets they have:						
	- Building Management System (BMS) with automated controls						
	- HVAC including chillers, air handling units, cooling towers						
	- Fans and pumps						
	- Other assets?						
5	If they have generators, then checking if:						
	- Do they know the size?						
	- Do they know how old are they?						
	- Are they tested regularly?						
	- Are they connected in parallel with the grid?						
	- Do they have a G59 connection?						
	- Other reasons?						
6	Discussing programme options						
7	Discussing potential impact to existing operations						
8	Offering free surveys, monitoring equipment and setup						

Question 1.4 - Additional Information – Initial Client Contact Closure Reasons

Please review the following closure reasons and:

1. Add any missing reasons or adjust any of the currently listed reasons.
2. Based on your experience please rank the closure reasons in the following table by likelihood of occurring based on a scale of 1 to 10 with 1 being the less likely reason and 10 being the most likely.

Initial Client Contact Closure Reasons	Likelihood of reason causing a client to be lost					
	Unlikely				Very Likely	N/A
	1	2	3	4	5	
Already with another DR provider						
No large assets or generators for DR						
Assets or generators deemed unsuitable for DR						
Not Interested - Impact Concern						
Not Interested or don't have enough time						
Not Interested - Insufficient Financial Returns						

Question 2.1 - Additional Information – Site Assessment Stage Steps

Please review the following stage steps and:

1. Add any missing steps or adjust any of the currently listed steps.
2. Based on your experience please indicate how important the task is when deciding if the site is suitable for Demand Response. With 1 meaning the task/outcome is not very important and 5 meaning this task/outcome is very important for deciding the suitability of the site.

Task	Site Assessment Stage Tasks	How important is this task in determining the suitability of a site for Demand Response?					
		Unimportant				Very Important	N/A
		1	2	3	4	5	
1	Obtain and analyse site information:						
	- Half hourly data obtained and analysed						
	- Assets list obtained and analysed						
	- If generator available then confirmation of size, status						
	- If generator available then confirmation of G59 status?						
2	If sufficient potential then gain initial agreement with client to continue assessment						
3	Delivery team confirms site potential and setup costs by either a site survey or phone survey						
4	Contract negotiations undertaken						
5	Project plan and framework created						
6	Once signed by client then the opportunity is turned into a Project and handed over to the Delivery Team						

Question 2.4 - Additional Information – Site Assessment Closure Reasons

Please review the following closure reasons and:

1. Add any missing reasons or adjust any of the currently listed reasons.
2. Based on your experience please rank the closure reasons in the following table by likelihood of occurring based on a scale of 1 to 10 with 1 being the less likely reason and 10 being the most likely.

Site Assessment Closure Reasons	Likelihood of reason causing a client to be lost					
	Unlikely				Very Likely	N/A
	1	2	3	4	5	
Out of time for current programme						
With another DR provider						
Lost interest - low priority						
Impact Concern						
Assets technically deemed unsuitable						
Upgrade costs						
Returns too low for KiWi Power						
Lost contact / Lost interest						
Returns too low for customer						
G59 issues (can't obtain or too expensive)						

Question 3.1 - Additional Information – Site Enablement and Go Live Stage Steps

Please review the following Site Enablement and Go Live stage steps and:

1. Add any missing steps or adjust any of the currently listed steps.
2. Based on your experience please indicate how important the task is when deciding if the site is suitable for Demand Response. With 1 meaning the task/outcome is not very important and 5 meaning this task/outcome is very important for deciding the suitability of the site.

Task	Site Enablement and Go Live Stage Tasks	How important is this task in determining the suitability of a site for Demand Response?					
		Unimportant			Very Important		N/A
		1	2	3	4	5	
1	Project Manager and Project Engineer assigned						
2	Handover performed by BDM						
3	Organise Technical Installation and Commissioning including:						
	- Arranging subcontractors						
	- Arranging equipment purchases						
	- Organising ENA applications						
	- PiP configuration						
	- G59 witness tests						
4	Training and User Acceptance Testing						
5	Spot testing						
6	Handover to Operations						

10 Appendix C

Appendix C contains the Likert scale results and minimum kW require results from the interviews given in chapter 3 for objective 1.

Table 39 - Chapter 3 Likert Scale Interview Results for Stage 1 Tasks

Task	Initial Client Contact Stage Tasks	Interviewee Results for: How important is this task in determining the suitability of a site for Demand Response? (1 = Low, 5 = High)														
		SS2	SJ1	SS1	SJ2	SJ3	SJ4	SI2	SI3	SI1	TS1	TS2	TJ1	Lower Q	Median	Upper Q
1	Prequalify potential site/client before making contact	4	4	2	2	4	3	2	2	4	N/A	N/A	N/A	2	3	4
2	Getting in contact with the right person in the company	4	4	4	2	5	4	4	4	5	N/A	N/A	N/A	4	4	4
3	Explaining Demand Response and checking to see if interested	4	4	4	4	3	4	1	3	4	N/A	N/A	N/A	3	4	4
4	Checking to see what turndown assets they have:															
	- Building Management System (BMS) with automated controls	3	2	5	3	4	4	2	5	5	N/A	N/A	N/A	3	4	5
	- HVAC including chillers, air handling units, cooling towers	5	3	5	5	4	4	5	5	5	N/A	N/A	N/A	4	5	5
	- Fans and pumps	5	3	5	5	4	4	5	5	5	N/A	N/A	N/A	4	5	5
	- Other assets?	5	4	5	5	4	4	5	5	5	N/A	N/A	N/A	4	5	5
5	If they have generators, then checking if:															
	- Do they know the size?	5	5	5	4	5	5	4	5	5	N/A	N/A	N/A	5	5	5
	- Do they know how old are they?	2	4	5	3	4	2	3	5	3	N/A	N/A	N/A	3	3	4
	- Are they tested regularly?	3	5	5	4	4	3	3	5	3	N/A	N/A	N/A	3	4	5
	- Are they connected in parallel with the grid?	4	5	5	4	5	4	5	5	3	N/A	N/A	N/A	4	5	5
	- Do they have a G59 connection?	4	5	5	3	5	4	5	5	3	N/A	N/A	N/A	4	5	5
	- Other reasons?	N/A	5	5	N/A	5	N/A	N/A	N/A	N/A	N/A	N/A	N/A	5	5	5
6	Discussing programme options	3	3	3	5	4	5	4	5	5	N/A	N/A	N/A	3	4	5
7	Discussing potential impact to existing operations	4	4	2	4	5	5	4	3	4	N/A	N/A	N/A	4	4	4
8	Offering free surveys, monitoring equipment and setup	2	3	1	4	5	3	4	3	4	N/A	N/A	N/A	3	3	4

Table 40 - Chapter 3 Likert Scale Interview Results for Stage 1 Reasons for Sites Being Lost

Initial Client Contact Closure Reasons	Interviewee Results for: Likelihood of reason causing a client to be lost? 1 = Low, 5 = High)														
	SS2	SJ1	SS1	SJ2	SJ3	SJ4	SI2	SI3	SI1	TS1	TS2	TJ1	Lower Q	Median	Upper Q
Already with another DR provider	2	3	3	2	3	2	3	1	3	N/A	N/A	N/A	2	3	3
No large assets or generators for DR	5	5	3	4	3	5	3	5	4	N/A	N/A	N/A	3	4	5
Assets or generators deemed unsuitable for DR	1	3	3	1	1	3	2	3	2	N/A	N/A	N/A	1	2	3
Not Interested - Impact Concern	4	5	3	4	4	5	4	2	4	N/A	N/A	N/A	4	4	4
Not Interested or don't have enough time	3	5	1	2	4	2	3	5	3	N/A	N/A	N/A	2	3	4
Not Interested - Insufficient Financial Returns	1	4	5	4	4	5	4	1	4	N/A	N/A	N/A	4	4	4

Table 41 - Chapter 3 Likert Scale Interview Results for Stage 2 Tasks

Task	Site Assessment Stage Tasks	Interviewee Results for: How important is this task in determining the suitability of a site for Demand Response? (1 = Low, 5 = High)														
		SS2	SJ1	SS1	SJ2	SJ3	SJ4	SI2	SI3	SI1	TS1	TS2	TJ1	Lower Q	Median	Upper Q
1	Obtain and analyse site information:															
	- Half hourly data obtained and analysed	3	4	5	5	5	5	4	5	5	N/A	N/A	N/A	4	5	5
	- Assets list obtained and analysed	5	4	5	4	5	5	5	5	5	N/A	N/A	N/A	5	5	5
	- If generator available then confirmation of size, status	5	4	5	5	4	5	5	5	5	N/A	N/A	N/A	5	5	5
	- If generator available then confirmation of G59 status?	5	4	5	2	4	5	4	3	3	N/A	N/A	N/A	3	4	5
2	If sufficient potential then gain initial agreement with client to continue assessment	4	5	5	2	5	3	4	3	4	N/A	N/A	N/A	3	4	5
3	Delivery team confirms site potential and setup costs by either a site survey or phone survey	4	5	5	3	5	5	4	5	5	N/A	N/A	N/A	4	5	5
4	Contract negotiations undertaken	4	5	1	1	4	5	5	5	5	N/A	N/A	N/A	4	5	5
5	Project plan and framework created	3	4	1	1	3	3	4	5	5	N/A	N/A	N/A	3	3	4
6	Once signed by client then the opportunity is turned into a Project and handed over to the Delivery Team	5	4	1	1	5	5	4	5	5	N/A	N/A	N/A	4	5	5

Table 42 - Chapter 3 Likert Scale Interview Results for Stage 2 Reasons for Sites Being Lost

Site Assessment Closure Reasons	Interviewee Results for: Likelihood of reason causing a client to be lost? (1 = Low, 5 = High)														
	SS2	SJ1	SS1	SJ2	SJ3	SJ4	SI2	SI3	SI1	TS1	TS2	TJ1	Lower Q	Median	Upper Q
Out of time for current programme	3	5	3	1	3	2	3	1	3	N/A	N/A	N/A	2	3	3
With another DR provider	2	3	3	2	2	4	3	1	2	N/A	N/A	N/A	2	2	3
Lost interest - low priority	3	4	2	2	2	2	1	4	4	N/A	N/A	N/A	2	2	4
Impact Concern	3	4	1	2	3	3	1	3	3	N/A	N/A	N/A	2	3	3
Assets technically deemed unsuitable	5	5	5	3	4	5	4	3	4	N/A	N/A	N/A	4	4	5
Upgrade costs	5	5	4	2	3	5	2	2	5	N/A	N/A	N/A	2	4	5
Returns too low for KiWi Power	4	3	4	3	4	2	3	3	4	N/A	N/A	N/A	3	3	4
Lost contact / Lost interest	3	3	1	1	4	1	4	1	3	N/A	N/A	N/A	1	3	3
Returns too low for customer	4	3	5	1	4	3	2	4	1	N/A	N/A	N/A	2	3	4

Table 43 - Chapter 3 Likert Scale Interview Results for Stage 3 Tasks

Task	Site Enablement and Go Live Stage Tasks	Interviewee Results for: How important is this task in determining the suitability of a site for Demand Response? (1 = Low, 5 = High)														
		SS2	SJ1	SS1	SJ2	SJ3	SJ4	SI2	SI3	SI1	TS1	TS2	TJ1	Lower Q	Median	Upper Q
1	Project Manager and Project Engineer assigned	4	3	N/A	N/A	4	4	3	5	5	2	5	3	3	4	5
2	Handover performed by BDM	5	3	N/A	N/A	5	4	4	5	3	3	4	3	3	4	5
3	Organise Technical Installation and Commissioning including:															
	- Arranging subcontractors	5	5	N/A	N/A	4	4	3	4	5	5	5	4	4	5	5
	- Arranging equipment purchases	4	5	N/A	N/A	4	4	4	4	5	3	4	4	4	4	4
	- Organising ENA applications	4	5	N/A	N/A	5	5	5	5	5	5	5	5	5	5	5
	- PiP configuration	4	5	N/A	N/A	5	5	5	5	5	1	5	5	5	5	5
	- G59 witness tests	3	5	N/A	N/A	5	5	5	5	5	5	4	5	5	5	5
4	Training and User Acceptance Testing	4	5	N/A	N/A	5	3	5	5	5	5	5	3	4	5	5
5	Spot testing	4	5	N/A	N/A	4	4	4	5	5	5	5	5	4	5	5
6	Handover to Operations	4	5	N/A	N/A	5	5	4	5	5	5	5	4	4	5	5

Table 44 - Minimum DSR kW Requirement Numbers for Figure 9

Interviewee	Min kW Asset Turndown	Min kW Turndown Site	Min kW Generator
SS2	50	200	500
SJ1	50	N/A	200
SS1	50	100	250
SJ2	80	250	500
SJ3	N/A	250	500
SJ4	200	200	500
SI2	100	200	250
SI3	200	200	400
SI1	500	500	1000
TS1	N/A	N/A	N/A
TS2	N/A	N/A	N/A
TJ1	100	N/A	750
Minamum	50	100	200
Lower Quartile	50	200	288
Median	100	200	500
Upper Quartile	200	250	500
Maximum	500	500	1000

11 Appendix D

Appendix D contains the detailed results from the DSR estimations in chapter 4 for objective 2.

Table 45 - MAPE and MBE Estimation Method Numbers for Figure 13

Method	MAPE			MBE		
	Average	Min	Max	Average	Min	Max
M1-V1	159%	96%	250%	150%	104%	236%
M1-V2	60%	35%	75%	-33%	-50%	1%
M2	56%	40%	70%	-16%	-41%	16%
M3	39%	33%	46%	-12%	-18%	-6%
M4	51%	39%	63%	-6%	-31%	5%

Table 46 - Sensitivity Analysis Result of Estimation Methods M1-V1, M1-V2 for Figure 14

(base values shaded in grey)

		Hotel 1										Hotel 2									
		2013					2016					2015					2016				
M1-V1	Input - Asset Usage %	40%	45%	50%	55%	60%	40%	45%	50%	55%	60%	40%	45%	50%	55%	60%	40%	45%	50%	55%	60%
	Input Difference % from Base	-20%	-10%	0%	10%	20%	-20%	-10%	0%	10%	20%	-20%	-10%	0%	10%	20%	-20%	-10%	0%	10%	20%
	Output - Total yearly MWh	1167	1313	1459	1604	1750	1168	1314	1460	1606	1752	1016	1143	1270	1397	1524	1017	1145	1272	1399	1526
	MWh Difference % from Base	-20%	-10%	0%	10%	20%	-20%	-10%	0%	10%	20%	-20%	-10%	0%	10%	20%	-20%	-10%	0%	10%	20%
	MAPE	139%	166%	193%	222%	251%	181%	215%	250%	284%	319%	80%	89%	98%	108%	118%	78%	86%	96%	108%	119%
	MAPE Difference % from Base	-28%	-14%	0%	15%	30%	-27%	-14%	0%	14%	28%	-19%	-10%	0%	10%	20%	-19%	-10%	0%	12%	24%
	MBE	78%	100%	122%	144%	167%	89%	112%	136%	160%	183%	169%	203%	236%	270%	303%	63%	84%	104%	125%	145%
M1-V2	Input - Asset Usage % of Base	5%	8%	10%	13%	15%	5%	8%	10%	13%	15%	5%	8%	10%	13%	15%	5%	8%	5%	8%	10%
	Input Difference % from Base	-50%	-25%	0%	25%	50%	-50%	-25%	0%	25%	50%	-50%	-25%	0%	25%	50%	-50%	-25%	-50%	-25%	0%
	Output - Total yearly MWh	178	268	357	446	535	153	230	307	383	460	192	287	383	479	575	190	285	178	268	357
	MWh Difference % from Base	-50%	-25%	0%	25%	50%	-50%	-25%	0%	25%	50%	-50%	-25%	0%	25%	50%	-50%	-25%	-50%	-25%	0%
	MAPE	65%	48%	35%	35%	42%	68%	63%	59%	58%	60%	85%	78%	71%	65%	63%	87%	81%	65%	48%	35%
	MAPE Difference % from Base	83%	37%	0%	-2%	19%	17%	7%	0%	-2%	2%	20%	10%	0%	-9%	-11%	17%	8%	83%	37%	0%
	MBE	-73%	-59%	-46%	-32%	-18%	-75%	-63%	-50%	-38%	-25%	-49%	-24%	1%	27%	52%	-69%	-54%	-73%	-59%	-46%
M1-V2	Input - Percentile %	3%	4%	5%	6%	7%	3%	4%	5%	6%	7%	3%	4%	5%	6%	7%	3%	3%	4%	5%	6%
	Input Difference % from Base	-40%	-20%	0%	20%	40%	-40%	-20%	0%	20%	40%	-40%	-20%	0%	20%	40%	-40%	-40%	-20%	0%	20%
	Output - Total yearly MWh	350	354	357	359	362	299	304	307	309	312	376	380	383	386	389	372	350	354	357	359
	MWh Difference % from Base	-2%	-1%	0%	1%	1%	-2%	-1%	0%	1%	2%	-2%	-1%	0%	1%	1%	-2%	-2%	-1%	0%	1%
	MAPE	36%	36%	35%	35%	35%	59%	59%	59%	59%	58%	71%	71%	71%	71%	71%	75%	36%	36%	35%	35%
	MAPE Difference % from Base	2%	1%	0%	0%	-1%	1%	0%	0%	0%	0%	1%	0%	0%	0%	-1%	1%	2%	1%	0%	0%
	MBE	-47%	-46%	-46%	-45%	-45%	-51%	-51%	-50%	-50%	-49%	-1%	0%	1%	2%	3%	-40%	-47%	-46%	-46%	-45%

Table 47 - Sensitivity Analysis Results of Estimation Method M2 for Figure 14

(base values shaded in grey)

		Hotel 1						Hotel 2					
		2013			2016			2015			2016		
M2	Input - Number of Clusters	3	4	5	3	4	5	3	4	5	3	4	5
	Input Difference % from Base	-25%	0%	25%	-25%	0%	25%	-25%	0%	25%	-25%	0%	25%
	Output - Total yearly MWh	349	403	391	429	442	344	417	440	499	390	551	539
	MWh Difference % from Base	-13%	0%	-3%	-3%	0%	-22%	-5%	0%	13%	-29%	0%	-2%
	MAPE	60%	57%	56%	60%	59%	61%	39%	40%	38%	69%	70%	65%
	MAPE Difference % from Base	7%	0%	-1%	1%	0%	4%	-2%	0%	-5%	-3%	0%	-7%
	MBE	-47%	-41%	-39%	-31%	-29%	-45%	10%	16%	32%	-38%	-12%	-14%

Table 48 - Sensitivity Analysis Results of Estimation Method M3 for Figure 14

(base values shaded in grey)

		Hotel 1								Hotel 2							
		2013				2016				2015				2016			
M3	Input - Number of Events	6	12	18	24	6	12	18	24	6	12	18	24	6	12	18	24
	Input Difference % from Base	-50%	0%	50%	100%	-50%	0%	50%	100%	-50%	0%	50%	100%	-50%	0%	50%	100%
	Output - Total yearly MWh	473	555	557	563	621	575	600	602	195	356	253	288	269	509	455	456
	MWh Difference % from Base	-15%	0%	0%	1%	8%	0%	4%	5%	-45%	0%	-29%	-19%	-47%	0%	-11%	-11%
	MAPE	41%	33%	32%	33%	41%	40%	40%	40%	64%	36%	53%	46%	69%	46%	48%	48%
	MAPE Difference % from Base	23%	0%	-2%	-1%	4%	0%	1%	0%	75%	0%	45%	28%	52%	0%	5%	5%
	MBE	-28%	-15%	-15%	-14%	0%	-7%	-3%	-3%	-48%	-6%	-33%	-24%	-57%	-18%	-27%	-27%

Table 49 - Sensitivity Analysis Results of Estimation Method M4 for Figure 14

(base values shaded in grey)

		Hotel 1										Hotel 2									
		2013					2016					2015					2016				
	Input - Set Point Value	25	24	23	22	21	25	24	23	22	21	25	24	23	22	21	25	24	23	22	21
	Input Difference % from Base	9%	4%	0%	-4%	-9%	9%	4%	0%	-4%	-9%	9%	4%	0%	-4%	-9%	9%	4%	0%	-4%	-9%
	Output - Total yearly MWh	554	598	647	690	741	551	598	650	697	751	304	342	384	434	491	331	384	431	484	544
M4	**MWh Difference % from Base**	-14%	-8%	0%	7%	14%	-15%	-8%	0%	7%	16%	-21%	-11%	0%	13%	28%	-23%	-11%	0%	12%	26%
	MAPE	53%	55%	58%	60%	63%	58%	60%	63%	66%	72%	48%	43%	39%	35%	33%	57%	50%	45%	40%	37%
	MAPE Difference % from Base	-8%	-5%	0%	4%	10%	-8%	-5%	0%	6%	14%	25%	12%	0%	-10%	-15%	27%	12%	0%	-10%	-18%
	MBE	-16%	-9%	-1%	5%	13%	-11%	-3%	5%	13%	21%	-20%	-10%	2%	15%	30%	-47%	-38%	-31%	-22%	-13%
	Input - Air Infiltration Value	0.5	0.6	0.7	0.8	0.9	0.5	0.6	0.7	0.8	0.9	0.5	0.6	0.7	0.8	0.9	0.5	0.6	0.7	0.8	0.9
	Input Difference % from Base	-29%	-14%	0%	14%	29%	-29%	-14%	0%	14%	29%	-29%	-14%	0%	14%	29%	-29%	-14%	0%	14%	29%
	Output - Total yearly MWh	802	772	647	516	381	823	780	650	501	343	480	463	384	298	221	521	506	431	344	267
M4	**MWh Difference % from Base**	24%	19%	0%	-20%	-41%	27%	20%	0%	-23%	-47%	25%	20%	0%	-22%	-43%	21%	17%	0%	-20%	-38%
	MAPE	76%	73%	58%	47%	52%	94%	84%	63%	51%	57%	36%	36%	39%	49%	60%	41%	41%	45%	55%	65%
	MAPE Difference % from Base	32%	26%	0%	-18%	-10%	49%	34%	0%	-18%	-10%	-8%	-8%	0%	26%	54%	-10%	-9%	0%	23%	45%
	MBE	22%	18%	-1%	-21%	-42%	33%	26%	5%	-19%	-44%	27%	23%	2%	-21%	-42%	-16%	-19%	-31%	-45%	-57%
	Input - U-Value Change %	-20	-10	0	10	20	-20	-10	0	10	20	-20	-10	0	10	20	-20	-10	0	10	20
	Input Difference % from Base	-20%	-10%	0%	10%	20%	-20%	-10%	0%	10%	20%	-20%	-10%	0%	10%	20%	-20%	-10%	0%	10%	20%
	Output - Total yearly MWh	692	669	647	627	608	691	670	650	631	613	422	402	384	370	357	471	450	431	415	401
M4	**MWh Difference % from Base**	7%	3%	0%	-3%	-6%	6%	3%	0%	-3%	-6%	10%	5%	0%	-4%	-7%	9%	4%	0%	-4%	-7%
	MAPE	61%	59%	58%	56%	55%	66%	65%	63%	61%	60%	36%	37%	39%	40%	42%	41%	43%	45%	47%	48%
	MAPE Difference % from Base	6%	3%	0%	-2%	-5%	6%	3%	0%	-2%	-4%	-8%	-4%	0%	4%	8%	-8%	-4%	0%	4%	7%
	MBE	5%	2%	-1%	-5%	-7%	12%	8%	5%	2%	-1%	12%	6%	2%	-2%	-6%	-24%	-28%	-31%	-33%	-36%

12 Appendix E

Appendix E contains the detailed results of the Load Profile based DSR estimation method as developed in chapter 5 for objective 3.

Table 50 - Data Selection Method Numbers for Figure 18

	Minimum	Q1	Median	Q3	Maximum	Variance
Out-Of-Sample MAPE	40.7%	51.9%	59.0%	61.5%	76.9%	7.9%
Out-Of-Sample MBE	-40.9%	-29.4%	-5.7%	-0.7%	9.6%	16.0%
K-fold Cross Validation Using Sites MAPE	49.6%	55.6%	63.3%	66.8%	69.8%	5.9%
K-fold Cross Validation Using Sites MBE	-7.8%	-5.1%	-4.2%	2.5%	5.9%	4.2%
K-fold Cross Validation Using Random Selection MAPE	45.0%	49.6%	60.5%	63.6%	67.8%	7.4%
K-fold Cross Validation Using Random Selection MBE	-10.5%	-8.7%	-6.3%	-0.1%	2.2%	4.2%

Table 51 - MAPE Versus Processing Time for Stochastic Sample Sizes Numbers for Figure 19

Stochastic Sample Size	Average MAPE	Processing Time in Minutes
100	59%	6.27
200	59%	6.45
300	59%	6.70
400	59%	6.95
500	59%	7.40
600	59%	7.72
700	59%	8.12
800	59%	8.48
900	59%	8.85
1000	59%	9.18
2000	59%	12.85
3000	59%	16.60
4000	59%	20.35
5000	59%	24.05
6000	59%	27.48
7000	59%	31.13
8000	59%	34.85
9000	59%	38.63
10000	59%	42.07

Table 52 - Clusters Input Values Numbers for Figure 20

Cluster	1	2	3	4	5	6	7	8	9	10
Average MAPE	63.6%	50.6%	49.3%	49.3%	49.2%	49.4%	49.2%	49.3%	49.2%	49.4%

Table 53 - Profile-Predictor Comparison Numbers for Figure 21

Cluster	1	2	3	4	5	6	7	8	9	10
Predictor - Day of Year	63.5%	63.7%	63.6%	63.5%	63.6%	63.6%	63.5%	63.6%	63.6%	63.6%
Predictor - Week of Year	63.5%	51.1%	49.8%	49.9%	49.9%	49.9%	49.9%	49.9%	49.8%	49.9%
Predictor - Month of Year	63.6%	50.6%	49.3%	49.3%	49.2%	49.4%	49.2%	49.3%	49.2%	49.4%
Predictor - Weekend of Weekday	63.5%	63.8%	63.8%	64.1%	63.6%	63.5%	63.6%	63.8%	64.1%	64.1%

Table 54 - Hotel HVAC Chiller Percentage Usage Boxplot Numbers for Figure 31

Month	Minimum	Lower Quartile	Median	Upper Quartile	Maximum
Jan	0.0%	1.8%	12.2%	16.8%	28.7%
Feb	0.0%	1.9%	12.3%	17.1%	28.9%
Mar	0.0%	3.2%	12.7%	17.0%	28.6%
Apr	0.0%	6.8%	14.2%	19.8%	32.9%
May	0.0%	14.9%	20.9%	28.4%	50.9%
Jun	6.2%	22.1%	31.5%	38.7%	51.8%
Jul	14.3%	28.9%	36.9%	44.3%	62.1%
Aug	15.0%	28.0%	34.8%	41.6%	59.1%
Sep	3.7%	19.1%	25.9%	35.8%	55.9%
Oct	0.0%	13.7%	19.1%	25.9%	45.2%
Nov	0.0%	6.5%	13.9%	19.2%	42.8%
Dec	0.0%	2.2%	12.2%	16.6%	28.4%
Average	3.3%	12.4%	20.6%	26.8%	42.9%

Table 55 - DSR Estimation Outcome Numbers for Figure 34

	MAPE			**MBE**		
Method	**Average**	**Min**	**Max**	**Average**	**Min**	**Max**
M1-V1	159.4%	96%	250%	149.6%	104%	236%
M1-V2	59.9%	35%	75%	-33.2%	-50%	1%
M2	56.4%	40%	70%	-16.1%	-41%	16%
M3	38.8%	33%	46%	-11.7%	-18%	-6%
M4	51.1%	39%	63%	-6.4%	-31%	5%
M5 - Upper CI	109.2%	81%	145%	107.6%	68%	175%
M5 - Upper Quartile	56.8%	39%	65%	33.6%	8%	78%
M5 - Median	46.5%	24%	65%	1.1%	-16%	38%
M5 - Lower Quartile	66.5%	64%	73%	-35.6%	-53%	-1%
M5 - Lower CI	71.1%	42%	92%	-83.6%	-87%	-76%

Table 56 - January to March Profile Numbers for Figure 22

January							February							March						
HH	97.5% CI	75% Upper Bound	50% Median	Daily Average	25% Lower Bound	2.5% CI	HH	97.5% CI	75% Upper Bound	50% Median	Daily Average	25% Lower Bound	2.5% CI	HH	97.5% CI	75% Upper Bound	50% Median	Daily Average	25% Lower Bound	2.5% CI
1	27%	16%	11%	12%	0%	0%	1	28%	16%	12%	12%	0%	0%	1	27%	17%	12%	13%	0%	0%
2	27%	16%	10%	12%	0%	0%	2	27%	16%	11%	12%	0%	0%	2	27%	16%	12%	13%	0%	0%
3	25%	16%	11%	12%	0%	0%	3	26%	16%	11%	12%	0%	0%	3	25%	16%	11%	13%	0%	0%
4	24%	16%	10%	12%	0%	0%	4	25%	16%	10%	12%	0%	0%	4	24%	16%	11%	13%	0%	0%
5	23%	15%	10%	12%	0%	0%	5	23%	15%	10%	12%	0%	0%	5	23%	15%	10%	13%	0%	0%
6	23%	15%	10%	12%	0%	0%	6	23%	15%	9%	12%	0%	0%	6	23%	15%	10%	13%	0%	0%
7	23%	15%	9%	12%	0%	0%	7	23%	15%	9%	12%	0%	0%	7	23%	15%	10%	13%	0%	0%
8	23%	15%	9%	12%	0%	0%	8	23%	15%	9%	12%	0%	0%	8	23%	15%	10%	13%	0%	0%
9	23%	15%	9%	12%	0%	0%	9	23%	15%	9%	12%	0%	0%	9	23%	15%	10%	13%	0%	0%
10	23%	15%	9%	12%	0%	0%	10	23%	15%	10%	12%	0%	0%	10	23%	15%	9%	13%	0%	0%
11	23%	15%	9%	12%	0%	0%	11	23%	15%	9%	12%	0%	0%	11	23%	15%	9%	13%	0%	0%
12	23%	15%	9%	12%	0%	0%	12	24%	15%	9%	12%	0%	0%	12	23%	15%	9%	13%	0%	0%
13	23%	15%	9%	12%	0%	0%	13	23%	15%	9%	12%	0%	0%	13	23%	15%	10%	13%	1%	0%
14	46%	16%	11%	12%	0%	0%	14	41%	16%	11%	12%	0%	0%	14	46%	16%	11%	13%	1%	0%
15	23%	16%	12%	12%	0%	0%	15	23%	16%	12%	12%	0%	0%	15	23%	16%	12%	13%	3%	0%
16	24%	16%	12%	12%	0%	0%	16	24%	16%	12%	12%	0%	0%	16	24%	16%	12%	13%	4%	0%
17	26%	17%	12%	12%	0%	0%	17	25%	17%	12%	12%	0%	0%	17	27%	17%	12%	13%	3%	0%
18	29%	16%	12%	12%	2%	0%	18	30%	17%	12%	12%	3%	0%	18	29%	17%	13%	13%	5%	0%
19	28%	17%	13%	12%	3%	0%	19	27%	17%	12%	12%	3%	0%	19	27%	17%	13%	13%	5%	0%
20	29%	17%	13%	12%	3%	0%	20	28%	17%	13%	12%	5%	0%	20	27%	17%	13%	13%	5%	0%
21	28%	17%	13%	12%	4%	0%	21	29%	17%	13%	12%	5%	0%	21	28%	17%	13%	13%	5%	0%
22	29%	17%	13%	12%	3%	0%	22	30%	18%	14%	12%	3%	0%	22	30%	18%	14%	13%	5%	0%
23	30%	18%	14%	12%	5%	0%	23	31%	18%	14%	12%	5%	0%	23	31%	18%	14%	13%	5%	0%
24	30%	18%	14%	12%	4%	0%	24	30%	18%	14%	12%	5%	0%	24	30%	18%	14%	13%	6%	0%
25	31%	18%	14%	12%	5%	0%	25	31%	18%	14%	12%	5%	0%	25	30%	18%	14%	13%	5%	0%
26	31%	18%	14%	12%	5%	0%	26	31%	18%	14%	12%	5%	0%	26	30%	18%	14%	13%	5%	0%
27	32%	18%	14%	12%	5%	0%	27	32%	18%	14%	12%	5%	0%	27	31%	18%	14%	13%	5%	0%
28	31%	18%	14%	12%	5%	0%	28	33%	18%	14%	12%	5%	0%	28	32%	18%	14%	13%	6%	0%
29	32%	18%	14%	12%	5%	0%	29	33%	19%	14%	12%	5%	0%	29	33%	19%	15%	13%	6%	0%
30	33%	18%	14%	12%	5%	0%	30	33%	19%	14%	12%	5%	0%	30	33%	18%	15%	13%	6%	0%
31	33%	19%	14%	12%	5%	0%	31	33%	19%	14%	12%	5%	0%	31	32%	19%	14%	13%	6%	0%
32	33%	18%	14%	12%	0%	0%	32	34%	19%	14%	12%	0%	0%	32	32%	19%	14%	13%	5%	0%
33	33%	18%	14%	12%	0%	0%	33	33%	19%	14%	12%	0%	0%	33	33%	19%	14%	13%	5%	0%
34	32%	19%	14%	12%	0%	0%	34	33%	19%	14%	12%	0%	0%	34	32%	19%	14%	13%	5%	0%
35	32%	18%	14%	12%	0%	0%	35	33%	18%	14%	12%	0%	0%	35	32%	18%	14%	13%	5%	0%
36	33%	18%	14%	12%	0%	0%	36	33%	19%	14%	12%	0%	0%	36	33%	19%	14%	13%	5%	0%
37	33%	18%	13%	12%	3%	0%	37	33%	19%	14%	12%	3%	0%	37	33%	18%	14%	13%	5%	0%
38	39%	19%	14%	12%	5%	0%	38	39%	19%	14%	12%	5%	0%	38	31%	19%	14%	13%	6%	0%
39	32%	18%	14%	12%	5%	0%	39	32%	18%	14%	12%	5%	0%	39	31%	18%	14%	13%	5%	0%
40	31%	18%	14%	12%	5%	0%	40	31%	18%	14%	12%	5%	0%	40	31%	18%	14%	13%	5%	0%
41	30%	17%	14%	12%	5%	0%	41	30%	18%	14%	12%	5%	0%	41	30%	18%	14%	13%	5%	0%
42	30%	18%	14%	12%	5%	0%	42	29%	18%	14%	12%	5%	0%	42	30%	17%	14%	13%	5%	0%
43	29%	17%	13%	12%	5%	0%	43	30%	18%	13%	12%	4%	0%	43	29%	18%	14%	13%	5%	0%
44	29%	17%	12%	12%	0%	0%	44	29%	17%	13%	12%	0%	0%	44	29%	17%	13%	13%	5%	0%
45	29%	17%	12%	12%	0%	0%	45	29%	17%	13%	12%	0%	0%	45	28%	17%	13%	13%	5%	0%
46	27%	17%	12%	12%	0%	0%	46	29%	17%	12%	12%	0%	0%	46	27%	17%	13%	13%	0%	0%
47	27%	16%	12%	12%	0%	0%	47	28%	17%	12%	12%	0%	0%	47	28%	17%	13%	13%	0%	0%
48	27%	16%	12%	12%	0%	0%	48	27%	16%	12%	12%	0%	0%	48	26%	17%	12%	13%	0%	0%

Table 57 - April to June Profile Numbers for Figure 22

April							May							June						
HH	97.5% CI	75% Upper Bound	50% Median	Daily Average	25% Lower Bound	2.5% CI	HH	97.5% CI	75% Upper Bound	50% Median	Daily Average	25% Lower Bound	2.5% CI	HH	97.5% CI	75% Upper Bound	50% Median	Daily Average	25% Lower Bound	2.5% CI
1	30%	19%	13%	14%	5%	0%	1	46%	26%	20%	21%	14%	0%	1	48%	36%	30%	32%	20%	3%
2	30%	19%	13%	14%	5%	0%	2	46%	26%	20%	21%	14%	0%	2	48%	35%	29%	32%	20%	0%
3	30%	18%	13%	14%	3%	0%	3	45%	25%	20%	21%	13%	0%	3	46%	34%	27%	32%	20%	0%
4	26%	18%	13%	14%	3%	0%	4	44%	25%	19%	21%	13%	0%	4	43%	34%	27%	32%	20%	0%
5	25%	18%	13%	14%	4%	0%	5	43%	23%	19%	21%	13%	0%	5	44%	33%	24%	32%	20%	0%
6	25%	18%	12%	14%	2%	0%	6	43%	23%	18%	21%	13%	0%	6	44%	33%	24%	32%	19%	1%
7	26%	17%	12%	14%	3%	0%	7	41%	23%	18%	21%	12%	0%	7	44%	33%	24%	32%	18%	0%
8	25%	17%	12%	14%	3%	0%	8	42%	23%	18%	21%	12%	0%	8	43%	33%	24%	32%	18%	0%
9	25%	17%	12%	14%	0%	0%	9	43%	23%	18%	21%	12%	0%	9	44%	33%	23%	32%	18%	0%
10	25%	17%	12%	14%	3%	0%	10	42%	23%	18%	21%	13%	0%	10	45%	33%	23%	32%	18%	0%
11	25%	17%	12%	14%	2%	0%	11	43%	23%	18%	21%	12%	0%	11	43%	33%	24%	32%	18%	1%
12	25%	17%	12%	14%	3%	0%	12	44%	23%	18%	21%	13%	0%	12	44%	33%	24%	32%	18%	0%
13	26%	17%	12%	14%	4%	0%	13	43%	23%	18%	21%	12%	0%	13	44%	35%	24%	32%	19%	4%
14	47%	18%	13%	14%	5%	0%	14	44%	23%	19%	21%	13%	0%	14	46%	34%	25%	32%	19%	5%
15	27%	18%	13%	14%	5%	0%	15	44%	23%	19%	21%	13%	0%	15	46%	35%	25%	32%	19%	4%
16	29%	18%	14%	14%	5%	0%	16	45%	24%	20%	21%	14%	0%	16	48%	36%	27%	32%	21%	5%
17	31%	19%	14%	14%	5%	0%	17	47%	26%	20%	21%	15%	0%	17	48%	36%	30%	32%	21%	5%
18	34%	20%	14%	14%	6%	0%	18	48%	27%	21%	21%	15%	0%	18	50%	36%	31%	32%	22%	6%
19	32%	19%	14%	14%	7%	0%	19	48%	27%	21%	21%	15%	0%	19	50%	37%	32%	32%	22%	8%
20	32%	20%	14%	14%	7%	0%	20	50%	29%	21%	21%	15%	0%	20	50%	38%	33%	32%	22%	6%
21	32%	20%	15%	14%	8%	0%	21	49%	29%	21%	21%	16%	0%	21	51%	39%	33%	32%	22%	5%
22	34%	21%	15%	14%	8%	0%	22	50%	30%	22%	21%	16%	0%	22	51%	40%	34%	32%	23%	6%
23	34%	20%	15%	14%	9%	0%	23	51%	31%	22%	21%	16%	0%	23	52%	41%	34%	32%	24%	10%
24	34%	21%	15%	14%	9%	0%	24	52%	31%	22%	21%	16%	0%	24	53%	41%	35%	32%	24%	9%
25	35%	21%	15%	14%	10%	0%	25	54%	32%	22%	21%	16%	0%	25	58%	42%	35%	32%	24%	9%
26	36%	21%	16%	14%	10%	0%	26	56%	32%	23%	21%	17%	0%	26	57%	42%	35%	32%	24%	10%
27	36%	22%	16%	14%	10%	0%	27	56%	32%	22%	21%	17%	0%	27	56%	42%	36%	32%	25%	9%
28	36%	22%	15%	14%	10%	0%	28	57%	33%	23%	21%	16%	0%	28	57%	44%	37%	32%	25%	11%
29	37%	22%	16%	14%	10%	0%	29	58%	33%	23%	21%	17%	0%	29	58%	43%	36%	32%	25%	11%
30	38%	22%	16%	14%	10%	0%	30	59%	33%	23%	21%	17%	0%	30	58%	44%	36%	32%	25%	12%
31	38%	22%	16%	14%	10%	0%	31	61%	33%	23%	21%	17%	0%	31	62%	44%	37%	32%	25%	11%
32	39%	22%	16%	14%	9%	0%	32	60%	33%	23%	21%	17%	0%	32	61%	45%	37%	32%	25%	11%
33	39%	22%	16%	14%	9%	0%	33	59%	33%	23%	21%	17%	0%	33	62%	44%	36%	32%	25%	9%
34	38%	22%	16%	14%	9%	0%	34	59%	34%	23%	21%	17%	0%	34	61%	44%	37%	32%	25%	9%
35	39%	22%	15%	14%	8%	0%	35	61%	33%	23%	21%	16%	0%	35	60%	44%	37%	32%	25%	8%
36	39%	22%	16%	14%	8%	0%	36	60%	33%	23%	21%	17%	0%	36	60%	44%	36%	32%	25%	9%
37	41%	22%	16%	14%	8%	0%	37	58%	33%	23%	21%	17%	0%	37	61%	45%	36%	32%	26%	6%
38	40%	22%	16%	14%	9%	0%	38	60%	33%	23%	21%	17%	0%	38	60%	44%	36%	32%	25%	10%
39	37%	21%	15%	14%	9%	0%	39	58%	32%	23%	21%	16%	0%	39	59%	43%	36%	32%	24%	10%
40	36%	20%	15%	14%	9%	0%	40	57%	32%	22%	21%	16%	0%	40	57%	42%	36%	32%	24%	6%
41	35%	21%	15%	14%	9%	0%	41	55%	31%	22%	21%	16%	0%	41	56%	42%	35%	32%	23%	9%
42	34%	20%	15%	14%	9%	0%	42	55%	31%	22%	21%	16%	0%	42	53%	41%	35%	32%	23%	10%
43	34%	20%	15%	14%	8%	0%	43	52%	30%	21%	21%	16%	0%	43	53%	40%	34%	32%	23%	11%
44	33%	20%	15%	14%	7%	0%	44	51%	29%	21%	21%	16%	0%	44	52%	39%	33%	32%	23%	9%
45	33%	20%	14%	14%	7%	0%	45	52%	29%	21%	21%	15%	0%	45	51%	39%	33%	32%	23%	8%
46	32%	19%	14%	14%	6%	0%	46	50%	28%	21%	21%	15%	0%	46	49%	37%	33%	32%	22%	8%
47	32%	19%	14%	14%	5%	0%	47	49%	27%	20%	21%	15%	0%	47	49%	37%	32%	32%	22%	6%
48	30%	19%	14%	14%	5%	0%	48	46%	27%	20%	21%	14%	0%	48	48%	36%	31%	32%	20%	5%

Table 58 - July to September Profile Numbers for Figure 23

July							August							September						
HH	97.5% CI	75% Upper Bound	50% Median	Daily Average	25% Lower Bound	2.5% CI	HH	97.5% CI	75% Upper Bound	50% Median	Daily Average	25% Lower Bound	2.5% CI	HH	97.5% CI	75% Upper Bound	50% Median	Daily Average	25% Lower Bound	2.5% CI
1	63%	41%	35%	37%	26%	10%	1	51%	38%	33%	35%	23%	12%	1	49%	33%	23%	26%	18%	0%
2	56%	40%	34%	37%	24%	12%	2	52%	38%	33%	35%	23%	12%	2	50%	33%	23%	26%	18%	0%
3	54%	39%	33%	37%	23%	12%	3	51%	36%	32%	35%	23%	12%	3	49%	32%	23%	26%	18%	0%
4	57%	38%	33%	37%	23%	8%	4	51%	36%	30%	35%	23%	9%	4	49%	32%	22%	26%	17%	0%
5	54%	38%	32%	37%	23%	10%	5	51%	36%	29%	35%	22%	10%	5	49%	31%	22%	26%	17%	0%
6	54%	38%	31%	37%	23%	10%	6	50%	36%	28%	35%	22%	12%	6	48%	29%	22%	26%	16%	0%
7	53%	37%	31%	37%	22%	10%	7	51%	35%	27%	35%	21%	9%	7	48%	28%	21%	26%	15%	0%
8	52%	37%	31%	37%	21%	10%	8	50%	36%	27%	35%	21%	12%	8	46%	29%	21%	26%	15%	0%
9	52%	37%	29%	37%	22%	11%	9	51%	35%	27%	35%	21%	10%	9	47%	27%	21%	26%	16%	0%
10	52%	37%	30%	37%	21%	10%	10	50%	35%	26%	35%	21%	10%	10	48%	27%	21%	26%	15%	0%
11	52%	38%	32%	37%	22%	12%	11	50%	36%	27%	35%	21%	10%	11	47%	29%	21%	26%	15%	0%
12	52%	37%	30%	37%	22%	10%	12	50%	35%	27%	35%	21%	9%	12	48%	29%	21%	26%	15%	0%
13	55%	39%	33%	37%	23%	10%	13	50%	37%	30%	35%	22%	10%	13	48%	32%	21%	26%	15%	0%
14	52%	39%	33%	37%	23%	12%	14	50%	36%	30%	35%	22%	11%	14	50%	33%	21%	26%	16%	1%
15	53%	39%	33%	37%	23%	12%	15	52%	36%	30%	35%	23%	13%	15	50%	33%	22%	26%	17%	0%
16	53%	40%	34%	37%	25%	13%	16	52%	38%	32%	35%	23%	14%	16	48%	33%	23%	26%	18%	3%
17	56%	41%	35%	37%	26%	15%	17	50%	39%	33%	35%	25%	15%	17	50%	34%	23%	26%	19%	1%
18	59%	42%	35%	37%	28%	15%	18	53%	39%	34%	35%	26%	16%	18	50%	34%	24%	26%	19%	5%
19	58%	43%	36%	37%	29%	14%	19	54%	41%	34%	35%	27%	15%	19	50%	36%	24%	26%	20%	5%
20	59%	44%	37%	37%	30%	15%	20	56%	41%	35%	35%	29%	15%	20	52%	36%	25%	26%	20%	6%
21	61%	44%	38%	37%	31%	15%	21	56%	42%	36%	35%	29%	16%	21	54%	37%	25%	26%	20%	5%
22	62%	45%	38%	37%	32%	15%	22	60%	43%	36%	35%	30%	17%	22	54%	37%	26%	26%	20%	5%
23	67%	47%	39%	37%	32%	15%	23	61%	44%	37%	35%	32%	17%	23	59%	38%	27%	26%	21%	6%
24	67%	47%	39%	37%	33%	16%	24	61%	44%	37%	35%	32%	17%	24	58%	38%	27%	26%	20%	6%
25	69%	48%	40%	37%	33%	17%	25	65%	45%	38%	35%	33%	18%	25	61%	39%	29%	26%	21%	6%
26	69%	48%	40%	37%	33%	17%	26	64%	46%	38%	35%	33%	18%	26	61%	39%	29%	26%	21%	7%
27	69%	49%	40%	37%	33%	16%	27	66%	47%	39%	35%	33%	18%	27	61%	40%	29%	26%	21%	7%
28	69%	49%	41%	37%	34%	17%	28	68%	47%	39%	35%	33%	18%	28	65%	40%	30%	26%	21%	8%
29	70%	50%	41%	37%	34%	17%	29	69%	47%	40%	35%	33%	18%	29	64%	41%	31%	26%	22%	7%
30	70%	50%	42%	37%	34%	18%	30	69%	48%	40%	35%	33%	18%	30	65%	41%	31%	26%	21%	9%
31	70%	50%	43%	37%	35%	17%	31	69%	48%	41%	35%	33%	18%	31	63%	42%	31%	26%	21%	9%
32	70%	50%	42%	37%	34%	18%	32	69%	48%	41%	35%	33%	19%	32	67%	42%	31%	26%	22%	6%
33	70%	50%	42%	37%	34%	17%	33	69%	48%	40%	35%	33%	20%	33	68%	41%	31%	26%	21%	5%
34	71%	51%	42%	37%	34%	17%	34	69%	48%	40%	35%	33%	17%	34	67%	42%	31%	26%	21%	6%
35	70%	51%	42%	37%	34%	17%	35	70%	48%	41%	35%	33%	18%	35	67%	41%	32%	26%	21%	5%
36	70%	50%	42%	37%	34%	17%	36	69%	47%	40%	35%	33%	18%	36	68%	42%	31%	26%	21%	5%
37	69%	50%	42%	37%	34%	16%	37	69%	47%	40%	35%	33%	17%	37	66%	41%	31%	26%	21%	5%
38	69%	50%	42%	37%	34%	17%	38	68%	47%	40%	35%	33%	18%	38	65%	41%	31%	26%	21%	6%
39	69%	50%	41%	37%	33%	16%	39	68%	47%	39%	35%	33%	17%	39	64%	39%	30%	26%	20%	6%
40	69%	49%	40%	37%	33%	15%	40	68%	46%	39%	35%	33%	17%	40	61%	39%	29%	26%	20%	6%
41	69%	48%	40%	37%	32%	15%	41	65%	45%	38%	35%	32%	17%	41	61%	39%	28%	26%	20%	7%
42	68%	48%	39%	37%	32%	16%	42	67%	44%	37%	35%	32%	18%	42	60%	38%	28%	26%	20%	8%
43	68%	47%	38%	37%	32%	16%	43	63%	44%	36%	35%	31%	18%	43	57%	37%	27%	26%	20%	7%
44	68%	47%	38%	37%	31%	16%	44	61%	42%	36%	35%	30%	17%	44	57%	37%	26%	26%	20%	5%
45	64%	45%	38%	37%	30%	16%	45	61%	42%	36%	35%	29%	18%	45	57%	36%	26%	26%	20%	5%
46	62%	44%	36%	37%	29%	14%	46	57%	41%	35%	35%	29%	16%	46	55%	35%	25%	26%	19%	0%
47	59%	43%	36%	37%	27%	14%	47	55%	40%	34%	35%	24%	15%	47	52%	34%	24%	26%	19%	3%
48	57%	42%	35%	37%	26%	13%	48	55%	38%	33%	35%	24%	13%	48	52%	34%	23%	26%	18%	0%

Table 59 - October to December Profile Numbers for Figure 23

October							November							December						
HH	97.5% CI	75% Upper Bound	50% Median	Daily Average	25% Lower Bound	2.5% CI	HH	97.5% CI	75% Upper Bound	50% Median	Daily Average	25% Lower Bound	2.5% CI	HH	97.5% CI	75% Upper Bound	50% Median	Daily Average	25% Lower Bound	2.5% CI
1	41%	24%	18%	19%	13%	0%	1	38%	18%	13%	14%	5%	0%	1	26%	16%	11%	12%	0%	0%
2	41%	24%	18%	19%	12%	0%	2	38%	18%	13%	14%	5%	0%	2	26%	16%	11%	12%	0%	0%
3	40%	23%	18%	19%	12%	0%	3	37%	18%	13%	14%	3%	0%	3	24%	16%	10%	12%	0%	0%
4	39%	23%	17%	19%	12%	0%	4	36%	18%	12%	14%	3%	0%	4	24%	15%	10%	12%	0%	0%
5	38%	22%	18%	19%	12%	0%	5	37%	18%	12%	14%	1%	0%	5	23%	15%	9%	12%	0%	0%
6	39%	22%	17%	19%	12%	0%	6	36%	17%	12%	14%	3%	0%	6	23%	15%	9%	12%	0%	0%
7	38%	21%	17%	19%	12%	0%	7	35%	17%	12%	14%	5%	0%	7	23%	14%	8%	12%	0%	0%
8	38%	21%	16%	19%	12%	0%	8	36%	17%	12%	14%	3%	0%	8	23%	14%	9%	12%	0%	0%
9	38%	21%	17%	19%	12%	0%	9	35%	17%	12%	14%	4%	0%	9	23%	15%	9%	12%	0%	0%
10	37%	21%	16%	19%	11%	0%	10	35%	17%	12%	14%	3%	0%	10	23%	14%	9%	12%	0%	0%
11	40%	21%	16%	19%	11%	0%	11	36%	17%	12%	14%	0%	0%	11	23%	14%	9%	12%	0%	0%
12	39%	21%	17%	19%	12%	0%	12	36%	17%	12%	14%	3%	0%	12	23%	15%	9%	12%	0%	0%
13	41%	21%	16%	19%	10%	0%	13	39%	17%	12%	14%	4%	0%	13	23%	15%	9%	12%	0%	0%
14	41%	22%	17%	19%	12%	0%	14	47%	18%	13%	14%	5%	0%	14	46%	16%	11%	12%	0%	0%
15	40%	22%	17%	19%	13%	0%	15	38%	18%	13%	14%	5%	0%	15	23%	15%	12%	12%	0%	0%
16	42%	23%	18%	19%	13%	0%	16	39%	18%	13%	14%	5%	0%	16	24%	16%	12%	12%	0%	0%
17	42%	23%	19%	19%	14%	0%	17	39%	18%	14%	14%	5%	0%	17	26%	16%	12%	12%	1%	0%
18	44%	24%	19%	19%	14%	0%	18	42%	19%	14%	14%	6%	0%	18	30%	17%	12%	12%	3%	0%
19	45%	24%	19%	19%	14%	0%	19	42%	19%	14%	14%	6%	0%	19	27%	17%	12%	12%	1%	0%
20	45%	25%	19%	19%	14%	0%	20	42%	19%	14%	14%	7%	0%	20	25%	17%	13%	12%	4%	0%
21	47%	25%	19%	19%	14%	0%	21	43%	19%	14%	14%	8%	0%	21	27%	17%	13%	12%	5%	0%
22	48%	27%	20%	19%	14%	0%	22	43%	20%	15%	14%	8%	0%	22	30%	17%	13%	12%	5%	0%
23	48%	26%	20%	19%	15%	0%	23	46%	20%	15%	14%	9%	0%	23	30%	17%	14%	12%	5%	0%
24	48%	27%	20%	19%	15%	0%	24	44%	20%	15%	14%	9%	0%	24	30%	18%	14%	12%	5%	0%
25	49%	29%	21%	19%	15%	0%	25	47%	21%	15%	14%	10%	0%	25	30%	18%	14%	12%	5%	0%
26	49%	29%	21%	19%	15%	0%	26	48%	21%	15%	14%	10%	0%	26	32%	18%	14%	12%	5%	0%
27	50%	29%	21%	19%	15%	0%	27	47%	20%	15%	14%	10%	0%	27	31%	18%	14%	12%	5%	0%
28	49%	30%	21%	19%	15%	0%	28	48%	21%	15%	14%	10%	0%	28	32%	18%	14%	12%	5%	0%
29	50%	30%	21%	19%	15%	0%	29	50%	21%	15%	14%	10%	0%	29	33%	18%	14%	12%	5%	0%
30	50%	30%	21%	19%	15%	0%	30	48%	21%	15%	14%	10%	0%	30	32%	18%	14%	12%	5%	0%
31	51%	31%	21%	19%	15%	0%	31	49%	21%	15%	14%	10%	0%	31	33%	18%	14%	12%	5%	0%
32	51%	30%	21%	19%	15%	0%	32	50%	21%	15%	14%	9%	0%	32	32%	18%	14%	12%	5%	0%
33	50%	31%	21%	19%	15%	0%	33	49%	22%	15%	14%	8%	0%	33	32%	18%	14%	12%	3%	0%
34	51%	31%	21%	19%	15%	0%	34	49%	21%	15%	14%	7%	0%	34	32%	18%	14%	12%	1%	0%
35	50%	30%	21%	19%	15%	0%	35	49%	21%	15%	14%	8%	0%	35	33%	18%	14%	12%	0%	0%
36	51%	31%	21%	19%	15%	0%	36	48%	21%	15%	14%	7%	0%	36	33%	18%	13%	12%	0%	0%
37	50%	31%	21%	19%	15%	0%	37	49%	21%	15%	14%	7%	0%	37	33%	18%	13%	12%	3%	0%
38	50%	31%	21%	19%	15%	0%	38	48%	22%	16%	14%	9%	0%	38	40%	19%	14%	12%	5%	0%
39	50%	30%	20%	19%	15%	0%	39	48%	21%	15%	14%	9%	0%	39	31%	18%	14%	12%	5%	0%
40	50%	29%	20%	19%	15%	0%	40	46%	20%	15%	14%	8%	0%	40	32%	18%	14%	12%	5%	0%
41	47%	28%	20%	19%	15%	0%	41	46%	20%	15%	14%	9%	0%	41	29%	17%	13%	12%	5%	0%
42	48%	28%	20%	19%	15%	0%	42	45%	20%	14%	14%	9%	0%	42	29%	17%	13%	12%	5%	0%
43	48%	26%	20%	19%	15%	0%	43	44%	20%	15%	14%	8%	0%	43	29%	17%	13%	12%	5%	0%
44	47%	25%	20%	19%	14%	0%	44	44%	19%	14%	14%	6%	0%	44	28%	17%	13%	12%	2%	0%
45	45%	25%	20%	19%	14%	0%	45	43%	19%	14%	14%	6%	0%	45	27%	17%	13%	12%	1%	0%
46	44%	25%	19%	19%	14%	0%	46	42%	19%	14%	14%	5%	0%	46	29%	17%	12%	12%	0%	0%
47	43%	24%	19%	19%	14%	0%	47	43%	19%	14%	14%	6%	0%	47	27%	16%	12%	12%	0%	0%
48	42%	23%	18%	19%	13%	0%	48	39%	18%	14%	14%	5%	0%	48	25%	16%	12%	12%	0%	0%

Table 60 - Comparison of Actual Usage Against Load Profile Numbers for Figure 24

HH	97.5% CI	75% Upper Bound	50% Median	Daily Average	25% Lower Bound	2.5% CI	Maida Vale Actual Usage 23/12/13
1	26%	16%	11%	12%	0%	0%	0%
2	26%	16%	11%	12%	0%	0%	0%
3	24%	16%	10%	12%	0%	0%	0%
4	24%	15%	10%	12%	0%	0%	0%
5	23%	15%	9%	12%	0%	0%	0%
6	23%	15%	9%	12%	0%	0%	0%
7	23%	14%	8%	12%	0%	0%	0%
8	23%	14%	9%	12%	0%	0%	0%
9	23%	15%	9%	12%	0%	0%	0%
10	23%	14%	9%	12%	0%	0%	0%
11	23%	14%	9%	12%	0%	0%	0%
12	23%	15%	9%	12%	0%	0%	0%
13	23%	15%	9%	12%	0%	0%	3%
14	46%	16%	11%	12%	0%	0%	51%
15	23%	15%	12%	12%	0%	0%	16%
16	24%	16%	12%	12%	0%	0%	13%
17	26%	16%	12%	12%	1%	0%	13%
18	30%	17%	12%	12%	3%	0%	13%
19	27%	17%	12%	12%	1%	0%	13%
20	25%	17%	13%	12%	4%	0%	13%
21	27%	17%	13%	12%	5%	0%	13%
22	30%	17%	13%	12%	5%	0%	13%
23	30%	17%	14%	12%	5%	0%	13%
24	30%	18%	14%	12%	5%	0%	12%
25	30%	18%	14%	12%	5%	0%	13%
26	32%	18%	14%	12%	5%	0%	12%
27	31%	18%	14%	12%	5%	0%	12%
28	32%	18%	14%	12%	5%	0%	12%
29	33%	18%	14%	12%	5%	0%	12%
30	32%	18%	14%	12%	5%	0%	12%
31	33%	18%	14%	12%	5%	0%	10%
32	32%	18%	14%	12%	5%	0%	0%
33	32%	18%	14%	12%	3%	0%	0%
34	32%	18%	14%	12%	1%	0%	0%
35	33%	18%	14%	12%	0%	0%	0%
36	33%	18%	13%	12%	0%	0%	0%
37	33%	18%	13%	12%	3%	0%	3%
38	40%	19%	14%	12%	5%	0%	41%
39	31%	18%	14%	12%	5%	0%	13%
40	32%	18%	14%	12%	5%	0%	13%
41	29%	17%	13%	12%	5%	0%	13%
42	29%	17%	13%	12%	5%	0%	12%
43	29%	17%	13%	12%	5%	0%	12%
44	28%	17%	13%	12%	2%	0%	13%
45	27%	17%	13%	12%	1%	0%	10%
46	29%	17%	12%	12%	0%	0%	0%
47	27%	16%	12%	12%	0%	0%	0%
48	25%	16%	12%	12%	0%	0%	0%