ERASMUS UNIVERSITEIT ROTTERDAM
ERASMUS SCHOOL OF ECONOMICS

Powered by the people: Assessing the effectiveness of the smart meter placement in the Netherlands.

Master Thesis

MSc Economics of Markets and Organisations

Erasmus School of Economics

Erasmus University Rotterdam

Author: M. D. Spee

Student Number: 431661

Supervisor: dr. B.S.Y. Crutzen

Second Assessor: dr. J. Delfgaauw

Final Version submitted: November 10, 2020

Acknowledgements

With the writing of my master thesis now coming to an end, I would like to take the opportunity to express my gratitude to several people for their contributions. First of all to my supervisor, dr. Benoit Crutzen, for the professional and creative insights he provided. He acted as a valuable source of inspiration when I needed it the most; I am sure he will become an even more successful professor and wish him a bright future. Second, I would like to thank my friend Harmco Dijkslag, who was willing to thoroughly read the analysis section. His useful comments on this part were of great help. Finally, a special thanks to my girlfriend Elise, who has been of great support during the writing process.

Abstract

In the past decade, the Dutch government enacted that at least 80 % of all small-scale electricity users should be equipped with a smart meter connection by the end of 2020. With this target date quickly approaching, it is time to assess the effectiveness of this smart meter implementation policy. In this thesis, we did so by selecting a group of postal code areas in which the proportion of smart meter connections substantially increased to 80 % on average in 2016. After that, we selected a group of identical areas in which no such a large-scale implementation of smart meters occurred. As a result, we were able to compare the developments in average electricity use between the two groups over time. With this reliable methodology, we found a statistically significant reduction in electricity consumption that is caused by the large-scale implementation of smart meters. However, the measured effect is much weaker than expected, hence questioning the effectiveness of the policy.

Keywords: *smart meter, electricity grid, energy savings, difference-in-difference analysis.*

Table of Contents

1 – Introduction

The extraction, production and use of energy sources have major environmental implications. Perhaps the most important implication is the emission of pollutants (United Nations Development Programme, 2000), such as carbon dioxide. During the 1960's and early 1970's, a truly global awareness on these environmental issues emerged (Haq & Paul, 2013; Lowe, 1985). At that time, new empirical knowledge and the – technical – possibilities to observe the devastating effects of human conduct across the globe, resulted in a wide recognition of the human ability to affect the natural environment. Although this increasing awareness has not specifically been a post-war development (Grove, 2002), it certainly revived and stimulated a new perception on the global environment and the sustainability of human conduct. Since then, many opportunities have been exploited to make energy consumption more sustainable, whilst other opportunities have only recently been discovered.

We can broadly divide these sustainability efforts into two categories. The first category comprises efforts that aim to reduce the ecological footprint at the supply side by extracting and/or generating energy sources more efficiently (e.g. Wüstenhagen & Bilharz, 2006). For instance, generating electricity by using windmills instead of coal facilities, which contributes to the reduction of air pollution. Efforts of the second category aim to reduce the total energy use at the demand side, mostly through altering behavior or through making devices and appliances more energy efficient (Behrangrad, 2015). In this research, we focus on one specific measure at the demand side which has become increasingly relevant in the past decade: the implementation of smart meters on large scale. Smart meters are advanced energy meters that measure the energy consumption of a consumer. But, compared to a regular energy meter, a smart meter also communicates real-time user data and provides added information to the utility companies (Depuru, Wang, & Devabhaktuni, 2011). This real-time communication entails several advantages, the most important one being the feasibility to provide feedback to the energy user. Providing feedback to consumers about their energy consumption stimulates more energy efficient behavior (e.g. Wilhite & Ling, 1995), which could potentially reduce energy consumption by up to 15 % (Darby, 2006).

In 2009, the European Commission introduced the Electricity Directive (Pöttering & Erlandsson, 2009). This directive intends to facilitate competition in the energy markets of the Member States, and to create a secure electricity infrastructure that allows for more efficient use of energy. One crucial precondition to achieve these objectives is the development of a smart grid, which is an advanced, automatic electricity delivery network that allows for unconventional power flows and the two-way exchange of information between consumer and net manager. (Zheng, Gao, & Lin, 2013). Therefore, through this directive, the European Commission demands the Member States to equip at least 80 % of their consumers – i.e. households, stores, offices etc. – with a smart electricity meter by the end of 2020. Prior to the implementation, however, each Member State had to economically assess all long-term costs and benefits to both the national energy market and the individual consumer. Whenever the expected costs of implementation for a specific country exceeded its projected benefits in the long run, the European Commission could exempt its government from the responsibility to install smart meters on large scale.

Already one year later, the Dutch government presented a positive assessment regarding the realization of a smart grid (Van Gerwen, Koenis, Schrijner, & Widdershoven, 2010). According to this study, the benefits of energy savings, lower call center expenses, market forces, and the redundancy of meter reading, outweigh the estimated costs and investments by € 770 million in terms of net present value. Of course, non-monetary benefits also play a role, since a smart network makes the energy grid more stable and better enables decentralized electricity generation through solar panels for example (Pöttering & Erlandsson, 2009). But, according to the Dutch assessment, the most important benefit is the direct energy savings caused by the feedback mechanism, amounting to € 1.47 billion. This benefit in monetary terms is primarily attributable to a real average decline of both 3.2 % in electricity use and 3.7 % in gas consumption.

The primary interest of our study is to find whether the provision of feedback on the use of electricity indeed induces a decline in electricity consumption. In general, feedback on energy use occurs either in a direct or in an indirect way (e.g. Darby, 2006). The former comprises any type of observations on the real-time consumption of the user, mostly through devices called In-Home Displays – IHD's. This is the most effective form of feedback in terms of the potential reduction in electricity use (Darby, 2006). The latter type contains user data that the energy provider obtains through the communication from the smart meters. Subsequently, the provider sends this information to the user in the form of an overview which is called the indicative overview of user costs. To guarantee the privacy of the user, the energy provider may obtain the user data only once every two months (Van Gerwen et al., 2010).

The main message of the economic assessment is that the provision of indirect feedback to the user should already be enough to induce a reduction of 3.2 % in electricity consumption. The assumptions thereby are that the proportion of smart meter connections equals at least 80 %, and that the percentage of users who switch off the functionality of the smart meter – and hence do not receive indirect feedback – is less than 2 %. However, during the course of the smart meter implementations, other reports already raised concerns that this type of feedback may not be as effective as the Dutch government supposed (Delmas, Fischlein, & Asensio, 2013; Schleich, Faure, & Klobasa, 2017). This makes it even more interesting to study the effectiveness of the smart meter implementation on large scale in the Netherlands.

1.1 – Research Objectives

Although we do not consider any heterogeneity in the provision of feedback to the consumer, our analysis is still valuable for several reasons. The first reason is academic in nature and entails a contribution to existing research that estimates the relationship between feedback and electricity consumption. Currently, this effect has only been studied in experimental settings, which leaves the societal impact of a large-scale placement of smart meters open to question. Hence, through examining whether this relationship also holds in the context of a large-scale implementation throughout the Netherlands, we aim to enhance current knowledge on the effect that information could have on the electricity consumption of an average user. Second, our findings also serve as a recommendation to other Member States of the European Union. As mentioned before, the obligation to equip at least 80 % of all consumers with a smart meter applies to all Member States, but only if such an implementation on large scale is beneficial from an economical perspective. Our findings could help to make a better assessment for governments that have been reluctant in their implementation efforts thus far. Specifically, if the large-scale rollout of a smart grid did not result in the expected reduction of electricity consumption. In that case, our findings can serve as a justification that allows governments to reconsider the economic assessment and revise the implementation policy as a result. Conversely, the implementation efforts in the Netherlands could serve as an example for other countries if we find an effect that corresponds to the expected reduction in electricity consumption.

Interestingly, despite the relevance of this topic following the introduction of the Electricity Directive, there has not been any scientific research that evaluated the actual effect in electricity use attributable to the placement of smart meters on large scale. Therefore, in this master thesis, we aim to find whether there is any effect, and – if so – to provide a reliable estimate of its magnitude. We do so by posing the following research question:

"What is the effect of a large-scale implementation of smart meters on the yearly average electricity consumption for a small-scale user in the Netherlands?"

This statement forms the guideline for our study. The expected reduction in electricity consumption by 3.2 % thereby serves as a hypothesis that we will empirically verify in the coming chapters.

1.2 – Thesis Outline

In order to provide an answer to our research question, we first have to understand the functioning of the Dutch energy market, and specifically what role a smart meter can play in this context. We explain these topics in Chapter 2, in which we describe how the Dutch energy market developed, and who the main players are. In addition, we elaborate on the functionality of the smart meter and we pay attention to existing literature on the effectiveness of a smart meter in the efforts to reduce electricity consumption. After that, we start our empirical analysis to examine whether the expected relationship is present. We do so by applying a quasi-experimental study on a panel dataset that contains all user data on energy consumption in the Netherlands over a time period from 2013 to 2018. In Chapter 3, we clarify the variables this dataset contains and we introduce the empirical methodology, which results in a specification of the hypothesized relationship between the variables used. After that, we continue to our analysis in the next chapter. We start this analysis by describing the selection procedure we apply to obtain our sample. After that, we show the most important sample characteristics and describe why this sample is suitable for our analysis. We conclude this fourth chapter by presenting the findings of the hypothesis test. We discuss these findings in Chapter 5, together with several implications and recommendations that follow from the results. Finally, we finish this master thesis with a conclusion in Chapter 6.

2 – Background

As we already noted in the Introduction, we first have to explore the context in which the smart meter implementations took place. Therefore, in this chapter, we describe the functioning of the Dutch electricity market and the main purposes of a smart meter. After that, we focus on the potential reduction in electricity consumption as the main aim of the implementation policy. Subsequently, we discuss scientific research that examined the relationship between information provision and electricity consumption behavior. At the end of the chapter, we mention in what way our research contributes to the existing knowledge.

2.1 – The Dutch Electricity Market

An electricity markets is a type of commodity market that deals with the generation and consumption of electrical energy (Lin, Magnago, Foruzan, & Albarracín-Sánchez, 2017). In a normal market, this would imply that parties supply electricity to other parties that consume electricity. In the 20th century, this was exactly how the electricity market functioned. In a specific region, one utility company generated electricity and maintained the power grid to serve the consumers within that area (Van Wezel & Van der Bie, 2015). This market organization of local monopolies and vertical integration was easy to understand, and even efficient from the perspective of the required infrastructure. However, the current electricity market is no such ordinary market anymore. Two major developments have caused the shift from these local monopolies towards a fully competitive market nowadays. First, the liberalization of the Dutch electricity market, which has been initiated by the introduction of the Dutch Electricity Act in 1998. This Electricity Act explicitly created a distinction between the suppliers of electric energy and the suppliers of the electricity network. Introducing competition for the latter type of suppliers would be very inefficient, because of the high expenses necessary to create and operate multiple networks in the same area. Therefore, the authority to operate the electricity network in the Netherlands has been geographically divided between net managers. Figure 1 shows the division of responsibilities for the network infrastructure between the seven net managers that are currently active in the Netherlands. To prevent any misuse of market power, the Electricity Act entitled the Dutch competition authorities to supervise the conduct of the net managers – and the entire electricity market in general (Van Wezel & Van der Bie, 2015).

FIGURE 1 - Regional Network Division between Net Managers

Hence, the introduction of competition has been primarily applicable to the suppliers of electricity – energy providers henceforth. These providers are allowed to use the entire electricity grid, such that they can serve any customer throughout the country. In this way, consumers can freely choose a provider that offers the most attractive tariff or bundle. However, to make matters more complicated, the providers do not generate the electricity they sell to the consumer. Instead, production of electricity takes place by independent producers, both on centralized level – i.e. through power plants – and decentralized level (Van Wezel & Van der Bie, 2015). These producers sell their generated electricity to so-called Program Responsible parties who subsequently sell the generated electricity to the energy providers and large-scale users. As already mentioned, the Electricity Act separated the task to operate the network from the production operations of the former utility companies. Thereby, the government also divided the responsibility for operating and maintaining the electricity grid in two categories. First, the transmission of electricity through the high voltage power lines – 110+ kV – which connect the power plants with the local infrastructure. An independent Transmission System Operator called TenneT operates this part of the network. The second category are the net managers, which we already identified as the parties responsible for the local electricity infrastructure. As we shall see, the net managers are of primary importance for our study, since the Dutch government made the net managers responsible for the large-scale implementation of smart meters.

The second important development that further accelerated the shift towards a fully competitive market is the fact that the distinction between suppliers and consumers of electricity has become less and less transparent. Before the introduction of the Electricity Act, as we explained, the positions and tasks of suppliers and consumers were clearly arranged. After 1998, the distinction between different types of suppliers already resulted in a new arrangement of responsibilities at the supply side of the electricity market. However, the obvious distinction between consumers and suppliers has become disturbed as well, especially because of the evolution in decentralized electricity generation in the past decade. This development caused the emergence of the so-called prosumers (e.g. Parag & Sovacool, 2016) who generate electric power for direct use or even store energy for later use (Zipperer et al., 2013). The problem of this relatively new development is the possibility that prosumers generate more electricity than they can use or store. The resulting Reverse Power Flows – RPF's – make administering the electricity grid much more complex (Sgouras, Bouhouras, Gkaidatzis, Doukas, & Labridis, 2017). On top of that, the increasing substitution of conventional energy sources by fluctuating renewable energy sources make the exact balancing of supply and demand an even more complicated challenge (Mattern, Staake, & Weiss, 2010). Whenever balancing supply and demand is not planned appropriately, RPF's can seriously damage the electricity grid (Sgouras et al., 2017). The increasing instability in the generation of electric power exposes the necessity for stabilizing mechanisms within the electricity grid. A smart grid – as defined in the Introduction – offers this stability, since it enables the network operators to keep better track of local generation, peaks in demand, capacity needs etc. (Zheng et al., 2013). Besides, a smart grid is also desirable from the perspective of a competitive market, because it allows consumers to switch more easily between different providers (Van Gerwen et al., 2010).

2.2 – The Functionality of Smart Meters

The fundamental instruments for a smart grid are smart meters, since these devices allow for a two-way exchange of information between net manager and consumer (Zheng et al., 2013). Therefore, in the context of a quickly developing electricity market, the Dutch government decided to make the electricity infrastructure ready for the future by transforming it to a smart grid (Van Wijnen, 2020). Earlier efforts started off with a legal proposal in 2008 that attempted to enact a countrywide implementation of smart meters (Vringer & Dassen, 2016). At that time, however, social resistance because of privacy- and security concerns caused a serious political discussion on the question whether a smart meter should be obligatory yes or no. As a result, the legislation has been adapted and enacted in 2012, with several amendments to ensure user privacy. Now, a central database stores all user data, which net managers and energy providers are allowed to access on particular occasions only. Moreover, an individual user can also choose to switch off or even refuse the smart meter device.

Of course, just like a regular electricity meter, a smart meter registers the electricity use of the consumer. But as we mentioned earlier, a smart meter has a supplementary feature: it can also communicate this consumption data. This is the main reason why people have been concerned about their personal privacy, since this communication implies that energy providers and net managers can observe the real-time consumption of a user. The current position that smart meters have in the Dutch smart grid prevent such infringements of privacy. To see why, we schematically depict the functionality of a typical smart meter in Figure 2. Note that, besides electricity use, the smart meter can also communicate supplementary data on gas- and water consumption. The P2 gate in Figure 2 facilitates this additional provision of information through connecting with the regular gas and water meters.

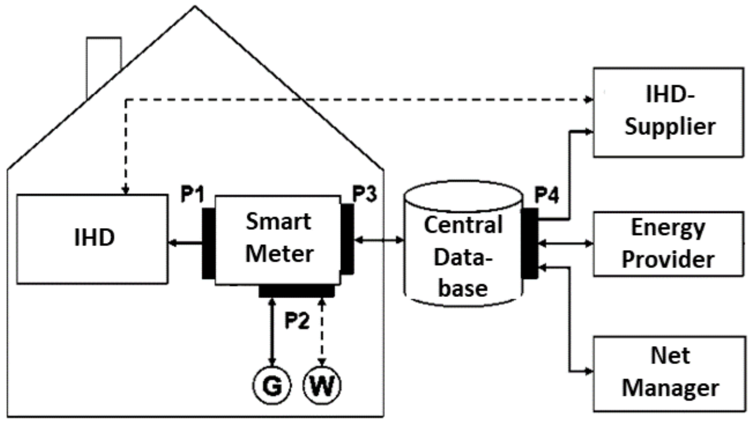

FIGURE 2 - Schematic Depiction of the Smart Meter Infrastructure Design

Essentially, a smart meter transfers information in two directions. The first information flow is directly towards the consumer – known as direct feedback – which passes through the P1 gate. Subsequently, an In-Home Display – IHD – receives this information and shows the real-time energy consumption to the user. In this case, the information stays within the residency or building. Hence, more important from a privacy perspective is the second type of communication: the flow of user data to the outside world. The Figure depicts this information flow from the smart meter, through the P3 gate, towards a central database. This database stores all information on the user's energy consumption, but is not immediately accessible to any other party. Privacy- and security rules restrict net managers and energy providers from having daily access to this central database. By default, user data is open to view once every two months for these parties, whereas more frequent access is allowed if and only if the consumer explicitly grants permission to retrieve their user data (Van Gerwen et al., 2010). If a consumer does not grant this permission, the P4 gate obstructs the service providers from retrieving the user data of this consumer. There are just three minor exceptions: at the end of the year, if a consumer moves away or switches from energy provider, the current provider has permission to compile a final electricity bill based on the stored user data. Ultimately, it is possible to switch off the smart meter whenever the design of operations still does not satisfy a consumer. In that case, the P3 gate is closed, and as a result, the user himself has to inform the energy provider about the annual electricity consumption.

The Dutch Smart Meter Requirements – DSMR – guarantee that all smart meters operate in the described way (Van Gerwen et al., 2010). Hence, through this standardization, the government assures a correct exchange of information between consumer and net manager, and protects consumers against any infringement of personal privacy. It thus becomes clear that consumers do not have to worry about any privacy issues related to smart meters. And if they do, they still have the option to refuse the installation of a smart meter.

But that raises another question: why should consumers accept one? How do they benefit from having a smart meter device? Van Gerwen et al. (2010). identify several advantages of having a smart meter over a conventional electricity meter. The first one is obvious: the automated communication implies that a consumer never has to send any user data anymore. Switching to another provider is thereby possible without any administrative efforts. Relatedly, the smart meter communication enables the energy providers to send a more accurate electricity bill which, in addition, can be specified with a lower electricity rate for off-peak consumption. The most important advantage, however, is that a smart meter can help consumers to gain a better understanding of their electricity use. More specifically, a smart meter provides the feasibility to provide feedback to the user, which raises more awareness among consumers on their electricity use. In the previous chapter, we already discovered two ways through which a smart meter can provide feedback. So far, we have mostly discussed direct feedback, which is the up-to-date communication about the current electricity use. Drawback of this type of feedback is that it requires an IHD, which can easily cost between € 150 and € 200,- (RVO, 2018). Obliging consumers to purchase such a display – despite its effectiveness (e.g. Burchell, Rettie, & Roberts, 2016; Darby, 2010) – is impossible, given the context of the designed legislation we described. However, one of the main societal purposes of the smart meter implementations has been to contribute to the reduction of electricity consumption (Vringer & Dassen, 2016). Hence, in order to comply with this purpose, the government intended to utilize the indirect type of feedback. To understand how this feedback mechanism works, we have to return to Figure 2 for a moment again. As we mentioned, this smart network design prevents net managers and other utility companies from obtaining confidential, real-time user data. However, with the consumer's permission, the Dutch government authorizes energy providers to access the central database once every two months in order to make providing feedback to the consumer possible. Based on the stored consumption data of each individual user over the past months, the energy provider creates a comprehensive overview of the electricity use and the related cost. This overview is called the indicative overview of user costs. The provision of information to the consumer through such an indicative electricity bill is a classic example of indirect feedback (Darby, 2006).

The Dutch Smart Meter Requirements prescribe that these indicative overviews must comply to several standards, such that all users receive feedback of equivalent quality. First, the indicative electricity bill should be clearly communicated to the user, meaning that passively distributing a user's indicative overview on a website portal is not acceptable. Second, it should contain both a historic comparison with the electricity use in previously recorded periods, and a normative comparison with the average consumption of a comparable user living in the same type of residency. Third, the overview should include an indication of the financial cost of the consumed electricity. Appendix A shows an example of an indicative overview.

2.3 – Previous Literature

According to estimations in the economic assessment of Van Gerwen et al. (2010), providing indirect feedback in this manner to all small-scale electricity users in the Netherlands should lead to a reduction of 3.2 % in the average electricity consumption of these users. In these estimations, Van Gerwen et al. (2010) and the Dutch government build on prior research by Darby (2006, 2010) that discusses the role of information provision on the behavior of consumers. In this section, we shortly review prior research on this relationship, thereby stressing the uniqueness of our research.

The concept of feedback creating consumer awareness has been an important dimension in the principles of psychology and communications for a long time (e.g. Bandura, 1969; Miller, Galanter, & Pribram, 1960). Presumably because of this strong theoretical base, the role of feedback became the focus of a variety of applied research studies as well (Hutton, Mauser, Filiatrault, & Ahtola, 1986). One of the areas in which the relationship between feedback and consumer behavior has been studied extensively, is the research on energy conservation. Already back in the 1980s, various empirical and conceptual analyses proved the usefulness of feedback for affecting energy consumption behavior (Cook & Berrenberg, 1981; Winett & Kagel, 1984). Ever since, scientists have contributed to this particular field of research on energy conservation. Several advancements that are worth mentioning are research on the effectiveness of different feedback types (Delmas et al., 2013; Fischer, 2008), feedback strategies (Abrahamse, Steg, Vlek, & Rothengatter, 2005; Cook & Berrenberg, 1981), environmental aspects that affect feedback effectiveness (e.g. Wood & Newborough, 2003), and the persistency of the relationship over time (Schleich et al., 2017).

The introduction of smart meter devices on the electricity market has added a new dimension to this field of research. As mentioned, a smart meter device provides several possibilities to provide consumption feedback to a user. Smart meters thereby allow for the generation of extremely detailed data on energy use. Together with the ongoing development towards a digitalized society, a much wider range of feedback possibilities has emerged. And more specific, digitalization and the abundance of data can be used to provide reliable, highly accurate feedback to each individual user on a large scale. As a result of these relatively new circumstances, we are now able to study the relationship between providing information and electricity use without using an experimental setting. In this way, our research provides a valuable contribution to the existing literature. As it appears, current research has almost exclusively been conducted in experimental settings (Delmas et al., 2013) or even through the use of surveys (e.g. Burchell, Rettie, & Roberts, 2016) to study the hypothesized effect. Correspondingly, by using accurate user data for almost all small-scale users in the Netherlands, we eliminate an important source of selection bias. And in this field of research, selection issues easily occur because of the many contingencies – as we just described – regarding the effectiveness of feedback on energy conservation behavior. This clearly emerges from earlier estimations of the effect of indirect feedback on average electricity use, which widely range from a 12 % reduction (Wilhite & Ling, 1995) to less than 2 % (Delmas et al., 2013).

3 – Concepts and Methodology

The problem statement and corresponding hypotheses defined in the previous chapters require that we obtain information on the use of electricity and how the smart meter network in the Netherlands unfolded over time. These are the central concepts of this study, with the former as the dependent variable and the latter as the independent variable. In this chapter, we explain these concepts and we show how these two variables evolved over time. Thereby we largely follow a protocol for data exploration to avoid estimating a misleading effect (Zuur, Ieno, & Elphick, 2010). The difficulty of testing our hypothesis is that many other variables likely influence the use of electric power. In Section 3 we provide an overview of all the factors that we can control for. However, despite the extensive number of control variables, we still find other important determinants which are impossible to control for. Therefore, we use a difference-in-difference analysis to minimize the possibility of any omitted variable bias. We describe this empirical methodology in the fourth Section, thereby specifying a mathematical expression of the expected relationship between the variables. We present the results of the hypothesis testing in Chapter 4.

3.1 – Dependent Variable

As mentioned, we use the consumption of electricity over time as the dependent variable, since we want to know whether the implementation of smart meters led to a reduction in electricity use. Unit of measurement for electricity use is the yearly average amount of kWh. Both the net managers as well as Statistics Netherlands – CBS – provide data on this consumption. However, the CBS provides estimates on household energy consumption only, whereas the data from the net managers comprises the consumption for small-scale users.

The latter one is more interesting for our research for the following reasons. First, the smart meter network is designed for all small-scale users, which also includes small shops and offices for example. Moreover, the intention of the policy implementation to reduce energy consumption by 3.2 % applies to all small-scale users, not only households (Van Gerwen, Koenis, Schrijner, & Widdershoven, 2010). Second, in areas that are not solely residential in nature, it is hard to differentiate between the household energy consumption and the consumption for non-domestic activities. This occurs, for example, in areas where a relatively high share of people works from a home office or live in a residency that is attached to their business. In such cases it is almost impossible to determine the true household electricity consumption. This allocation problem means that if we use household electricity consumption for our analysis, we cannot make a precise estimation of the hypothesized effect. The third reason is that using the consumption of all small-scale users gives us several convenient advantages in conducting our analysis. For instance, it allows us to control for any technical aspects that affect the electricity consumption, such as the type of connections. The technical information – provided by the net managers as well – applies to all small-scale connections per postal code area, and not just for the households. We further discuss the definition of small-scale users and the relevance of the technical factors in Section 3.

Trend in Energy Consumption

Hence, we use the average amount of kWh per active connection within a postal code area as the measure for electricity consumption. Formally, this consumption is the Standardized Yearly Consumption – SJV – which is the average consumption adjusted for local differences in weather conditions and normalized for exceptional circumstances such as power outages (Liander, 2020; Stedin, 2020). These adjustments are necessary to make appropriate comparisons of consumption possible, since even small differences in temperature and/or hours of sunshine can cause changes in space heating behavior and hence electricity consumption (Hart & De Dear, 2004). On top of that, it comprises the total consumption that is delivered through the grid and generated by own sources such as solar panels. We further discuss the role of own electricity generation in Section 3.

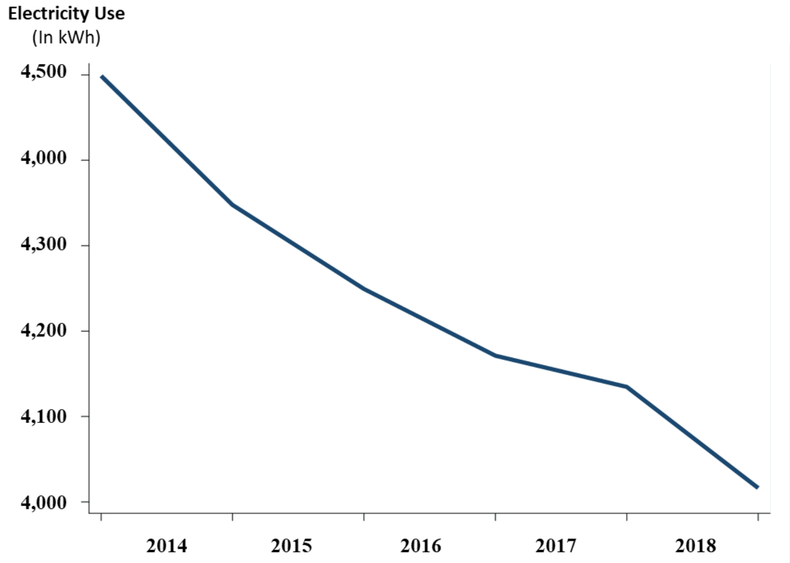

FIGURE 3 - Average Electricity Use of Small-scale Users in the Netherlands over Time

Overall, this average consumption equaled 4,235 kWh during the period, which is substantially higher than the average electricity consumption of 2,955 kWh of a Dutch household over the same period. This difference is simply the result of also including other small-scale users – shops, small offices – that appear to consume more energy on average. Interestingly, over the period, the per connection average shows a steady decline. Figure 3 shows how the average consumption decreased with 10.7% between 2013 and 2018, which is comparable to the 11.4% decrease in household electricity consumption in the Netherlands over the same period (CBS, 2020). This may indicate that, within the group of small-scale users, households have a relatively high potential for effectively reducing total electricity consumption (e.g. Yohanis, 2012). Two important developments induce this presumption: the improving efficiency of electronic appliances, and the rising consumer awareness of more energy-efficient household devices over the past decade (Pothitou, Hanna, & Chalvatzis, 2016). One development that can contradict this presumption is that the reduction in average electricity use is simply the result of a decrease in average household size (e.g. Kaza, 2010). However, the observations on average household size in our dataset show that this decrease – from 2.19 residents in 2013 to 2.15 by the end of 2018 – is negligible.

Our dataset also includes another important factor that may explain the decline in average electricity use. Over the period, 161,160 new dwellings have been constructed. Relatedly, the average energy label of the residencies improved. Postal code areas with a relatively high number of newly constructed houses also show a substantial improvement in the average value of energy labels within that area. Regarding other small-scale users, we see a similar interaction between energy label values and year of construction. Although we do not empirically test this interaction, it provides a reasonable indication that newly constructed buildings are more energy-efficient and hence contribute to the explanation of a declining average consumption.

Data Distribution

We continue to analyze the dependent variable by exploring the distribution of observations over the time period. Appendix B shows the yearly distributions in the observations on average electricity use per area. Clearly visible are two outliers in 2015 and 2016, where the average consumption exceeds 50,000 kWh. Both observations come from the same postal code, an industrial area in Etten-Leur, Noord-Brabant, with just ten inhabitants populating the area. During the other years the electricity consumption in this area was also high, but never exceeded 25,000 kWh on average. We can therefore confidently identify the observations for 2015 and 2016 as measurement errors and correspondingly remove them from our dataset (Zuur et al., 2010).

The variance between the observations seems to be homogeneous over the years. In each year, there is hardly any variation between observations in the first quartile of observations, and the spread in the upper quartile arises from no more than 63 postal code areas – i.e. 1.6% of the yearly observations – with an average consumption that exceeds 15,000 kWh. These areas are mostly industrial in nature. Typically, in such areas, the number of businesses per inhabitant is relatively high, together with a high percentage of heavy connections and on average 90% of the inhabitants employed. Contrarily, the first quartile of observations is mostly on densely populated postal code areas with a high degree of urbanization and a large share of Multi-Dwelling Units, such as apartments. Apart from two postal code areas, the average consumption for these observations does not fall below 2,000 kWh. A high share of unoccupied buildings in these two areas probably explains this low level of consumption.

To summarize, the data shows a right-skewed distribution of observations on the average energy consumption. This skewness persists over the years throughout the time period, and therefore also explains why the average consumption of all small-scale users is considerably higher than the household electricity consumption.

3.2 – Independent Variable

With this study, we aim to find whether there is a correlation between providing feedback to the consumer through smart meters on the one hand, and the average electricity use over time on the other hand. Hence, these smart meters serve as the mechanism that provides information to the consumers, which makes them the independent variable in our analysis. The measure we use is the percentage of the connections within a postal code area that have a smart meter device installed.

At the start of 2013, the percentage of smart meter connections was 5.3 %. Presumably, these smart meters have been installed before the large-scale implementation started – as described in the previous chapter. Figure 4 shows a constant annual increase of roughly five percent point during the first two years of the period, which results in an average implementation rate of 14.7 % by the end of 2014. A steady replacement of defective or outdated conventional meters likely explains this upward trend. After the policy has been enacted in 2015, we discover a more rapid increase. By the end of 2018, on average 61.4 % of the total connections within a postal code area had a smart meter implemented. Appendix C shows the distribution of the percentage of smart meters installed per postal code area over the years. Clearly visible is the effectiveness of the policy for the treated areas; from 2016 onwards, a new bell-shaped curve emerges with a median of approximately 80 % smart meter connections.

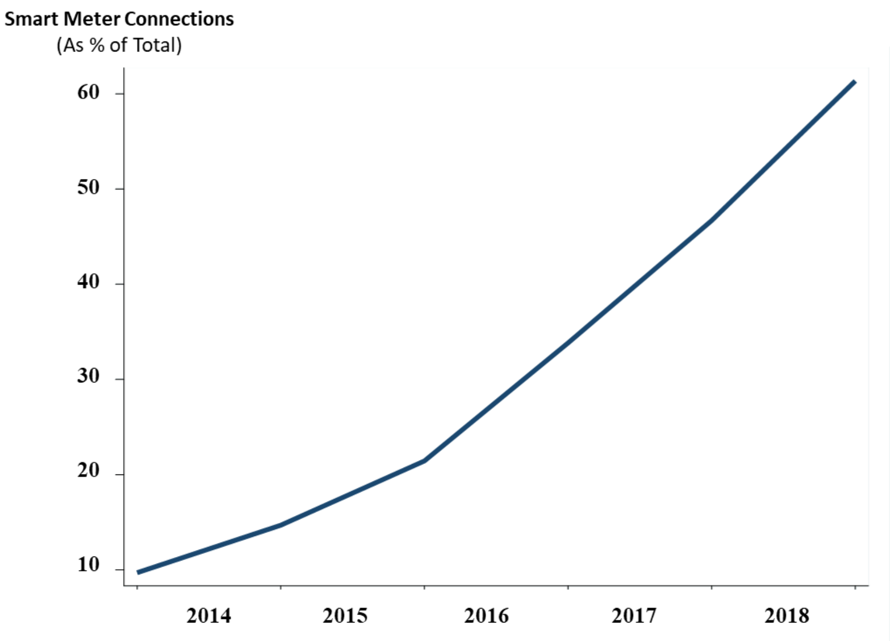

FIGURE 4 - Average Percentage of Smart Meter Connections over Time

When we further analyze the spread of smart meter devices, we find that some of the observations show inconsistencies on the percentage of smart meters over time. After we subtract the percentage of smart meter connections at the beginning of the year from the percentage at the end of the year – i.e. the implementation rate for a specific year – we find 82 observations where the implementation rate is negative. This can either be the result of the so-called smart meter refusals, or simply because of measurement errors. A refusal occurs if a consumer fully rejects the installation of a smart meter installation in the first place, but can also partly occur in case the user switches off the functionality of the smart meter. In the latter case, the smart meter does not communicate information on the energy consumption anymore, and consequently the user does not receive the bimonthly overview of user cost. Unfortunately, the net managers do not provide any information whether the measured percentage of smart meter connections includes or excludes the smart meters that have been switched off. As a result, we cannot identify if the negative implementation rates for 77 unique postal code areas are the result of measurement errors. Hence, we do not remove these observations from our dataset.

3.3 – Control Variables

What becomes clear in the first section of this chapter is that there are many factors that affect the use of electric energy. In order to find an unbiased estimator of the effect of smart meters, we have to take into account as many factors as possible. In this section, we provide an overview of all the factors we control for in our analysis. Thereby we classify these control variables in four categories: consumption data, technical details, controls on postal code level and population controls.

Consumption data

Since we use the Standardized Yearly Consumption to measure electricity consumption, controlling for weather circumstances and power outages is not needed. What is necessary, however, is to control for the use of alternative sources of energy, which are gas consumption and district heating. Electricity consumption may act as a substitute for the use of gas in cooking – by means of glass-ceramic or induction stoves – and heating. Hence, controlling for gas consumption prevents us from overestimating the hypothesized effect. However, the effect of controlling for district heating is uncertain. On the one hand, district heating can certainly help to reduce energy consumption for heating purposes, but a connection to the heating network often requires a heat pump that also uses a considerable amount of electricity.

An important aspect of energy consumption is that consumers have several possibilities to contribute to their own electricity supply, for example by means of solar panels. Our dependent variable already includes such own generation of electricity, often referred to as net-metering (Gillingham, Deng, et al., 2016). However, it is still important to include this aspect since it may reveal something about peoples' attitudes towards energy conservation behavior (Gillingham, Rapson, & Wagner, 2016).

Technical details

In the first section of this chapter we explained why we use the consumption data on all small-scale users instead of household consumption only. Small-scale users are defined by the maximum capacity of their connection (Liander, 2020; Stedin, 2020), i.e. the maximum number of hourly kilowatts that can flow through the connection. In the Netherlands, twelve types of small-scale connections exist, which range from 1x10 – i.e. the lowest capacity, mostly used for garage boxes – to 3x80. The first digit refers to the number of fuses, and the latter two digits refer to the maximum current that can flow through the connection. The most common connections are 1x35 and 3x25, which are the small- and standard household connections, respectively. The alternative for a small-scale connection is a large-scale business connection, which is heavier than the 3x80 connection and specifically designed for medium-sized and large businesses only. We control for differences in the connection distribution between postal code areas by creating dummy variables that measure the percentage of each specific type of connection as part of the total connections within an area. Since 1x10 and 1x20 connections are so uncommon, we combine these connection types together with the 1x25 connections in the lowest category. Other technical factors that may affect the consumption of electricity are the total length of the cabling and the average transformer capacity in the postal code area. These factors are both associated with the transportation efficiency of electricity, which decreases with a longer cabling distance and/or a higher transformer capacity (Georgilakis, 2011). This in turn can lead to a slight overestimation of the true electricity use, which is why we control for these factors.

Despite the three advantages of using the small-scale user data on electricity use we mentioned in the first section, there is also one obvious limitation: we cannot make a distinction between the other types of small-scale users. As a result, we cannot control for the type of buildings in which non-residential energy consumption takes place. In the next chapter, we provide more detail on how we take this issue into account.

District controls

The third category consists of factors that concern the urban characteristics of the postal code area. These characteristics include the address density and degree of urbanization of the area, together with various specifications on the residencies. Non-residential factors that we also include are the number of businesses, the total of electric vehicles registered, and the number of charging stations designed to recharge these vehicles. Regarding the housing characteristics, Kaza (2010) points out that the type of residency likely has an impact on the use of electricity. We categorize these types into two groups, Single-Dwelling Units – SDU's – and Multi-Dwelling Units – MDU's. The former includes all detached and semi-detached houses, whereas the latter includes any type of multi-family homes such as apartments and townhouses. We also include the average living space for both groups, because differences in floor area strongly affect the average electricity consumption across different types of residencies (Yohanis, Mondol, Wright, & Norton, 2008). Thereby, the average living space in terms of square meters is typically lower for MDU's compared to SDU's.

Another important housing characteristic which we already discussed is the year of construction. Aksoezen, Daniel, Hassler, & Kohler (2015) provide evidence that newly built houses are typically more energy efficient, after controlling for differences in consuming behavior. Interestingly, these authors show that there is no linear relationship between the year of construction and the consumption of electricity per cubic meter of housing space. Therefore, we classify the residential buildings into eight construction categories, with the first category containing the residencies constructed before 1945 and the final category the residencies constructed after 2014. However, as it appears, the relationship between electricity consumption and year of construction can also be misleading. Obviously, because the energy performance of a building can also be improved over time, for example because of renovations or the installation of more efficient electrical equipment. Controlling for all these individual improvements is impossible, but we can approximate such efforts by using the average energy label within the postal code area. An energy label shows the efficiency performance of a building and the corresponding possibilities to realize additional energy savings. The lower the label score, the more efficient the residency is. Hence, a reduction in the average label score over time indicates a general improvement of the energy efficiency within the postal code area. Of course, such an approximation can also be the result of self-selection – people with less energy efficient houses likely feel embarrassed and hence do not want to apply for an energy label – or the fact that newly constructed buildings are more energy efficient. Although the construction categories easily allow us to control for the latter, the self-selection effect is hard to identify. Still, it can help us to identify and compare possible trends of improving energy efficiencies within a postal code area over time, especially if those areas are very similar to each other. In order to control for other small-scale users as well, we include the average energy label per year for non-residential buildings in the same area. To further complement this approximation, we also include the number of subsidies granted for renovations that improve energy performance, together with their total amounts within a postal code area.

Other evident factors that are necessary to control for is the average household size and the share of one-person households (e.g. Kaza, 2010), house ownership (e.g. Ndiaye & Gabriel, 2011) and average occupancy rate.

Population controls

The final category includes the characteristics of the residents that live within the postal code area. According to previous research, the most important characteristics that we need to include are average age and the number of children (Yohanis et al.), income – both on individual and household level – (e.g. Cayla, Maizi, & Marchand, 2011), gender and ethnical background (Brounen, Kok, & Quigley, 2012), and employment status (Pothitou et al., 2016) Apart from average age and income, we measure these variables as a percentage of the total population within a specific area.

Before we continue to describe our analysis, we notice that one obvious control is missing in our description: energy prices (e.g. Asafu-Adjaye, 2000; Reiss & White, 2008). In the Netherlands, the price for electricity comprises taxes, a fixed component – a connection fee plus rent for the meter installed – and a variable component. The latter is the cost of delivery in terms of euro per kWh, which depends on the type of contract the consumer has with the energy provider, reductions during off-peak hours and whether the user also consumes gas. As a result, the actual price per kWh consumed can vary substantially between consumers (ACM, 2019). What further complicates a possible control for energy prices is the fact that an increasing number of small-scale users – 1.3 million users over 2018 – regularly switches from energy provider (ACM, 2019).

Hence, it becomes clear that we cannot provide a reliable answer to our research question if we apply a standard regression analysis. There are too many factors that we cannot control for, such as differences in energy contracts or electricity prices, renovations that improve energy performances, differences in the types of small-scale users across different areas, etc. On top of that, as Nielsen (1993) shows, factors related to lifestyle – e.g. the number of weeks on holiday, the use of household appliances – are at least half as important as the socioeconomic factors we mentioned in this section. Therefore, we take a different approach that allows us to confidently assume that electricity prices, responsiveness to prices, individual lifestyles and any other unobserved factor do not substantially differ between the selected postal code areas over the time period. In the next chapter, we elaborate more on this approach and why it justifies this assumption.

3.4 – Data and Analysis

In order to carry out our research, we first construct a panel dataset which allows us to track the mentioned variables over a time period from 2013 until 2018. This time period is specifically interesting, since it allows us to compare postal code areas over several consecutive year before the moment the policy was adopted. Unfortunately, by the time we conducted this analysis, there was still no public data available on housing and income in 2019. Therefore, we exclude all observations for 2019 from our analysis.

As mentioned, the policy has been introduced with the purpose to make at least 80% of the small-scale connections having a smart meter in 2020. The three large net managers in the Netherlands – Enexis, Liander and Stedin – publish yearly user data on connections and consumption within a postal code area over this period. The user data contains observations on 3,885 distinct postal codes at the four-digit level, hence covering 95% of all the postal code areas in the Netherlands. The missing 187 areas are mostly located in Zeeland, a province where Enexis, Liander and Stedin do not serve the market.

Since we need many control variables to accurately analyze the changes in electricity consumption, we also make use of two other sources of data. The first source is the regional statistics dataset from the

Statistics Netherlands (CBS, 2020). Statistics Netherlands publishes public data on demographic factors, housing and income on a yearly basis, both at four-digit and six-digit postal code level. We use the four-digit level, because at the latter level, many observations are omitted due to privacy concerns. The second source is Klimaatmonitor, which is a platform managed by Rijkswaterstaat (Rijkswaterstaat, 2020). Klimaatmonitor connects data sources from several government institutions together to one database that comprises many of the efforts that address environmental issues in the Netherlands. From these institutions, the Netherlands Enterprise Agency, the Netherlands Vehicle Authority, and Eco-Movement are of particular interest for our study, since they provide data on important control variables such as renovation subsidies, energy labels and the number of electrical vehicles registered per postal code area.

The entire dataset that results from connecting the different data sources contains a total of 22,513 observations on 3,885 different postal code areas. These postal code areas form the unit of analysis. As our dataset contains cross-sectional observations on each four-digit postal code over time, we have to apply a panel data study. Thereby we deliberately do not follow a standard fixed effects model analysis to explain the variation in average electricity consumption within each area (e.g. Stock & Watson, 2010). Instead, we use a more advanced methodology, which is necessary in order to estimate such a small average effect of 3.2 %. We select two groups of postal code areas; one in which a large proportion of smart meters had been implemented during a specific year – i.e. the treatment group – whereas in the control group no such implementation occurred. After including all postal codes that satisfy the implementation criteria, we can compare the change in average electricity consumption per postal code area between the two groups. More specifically, after 2015 we should find a stronger decrease in consumption in the treatment group compared to the control group.

This is what the difference-in-difference estimator does. Mathematically, we can express this relationship in the following equation:

$$C_{i,t} = \alpha_0 + \gamma_0\, PC_i + ß_1\, D_i + ß_2\, D_i\, T_t + \gamma_t + X_{i,t}\,\Gamma + \varepsilon_{i,t} \qquad (1)$$

where D_i is a binary variable that equals one if postal code area *i* belongs to the treatment group, and T_t a binary variable that equals one in year *t* after the implementation of the policy. In this expression, the $ß_2$ estimator equals the relationship of interest between the smart meter implementation and the average electricity consumption C for area *i* in year *t*. Furthermore, we add a vector of variables $X_{i,t}\,\Gamma$ to control for any possible differences between the treatment and control group. These variables comprise the four control categories that we described in Section 3. Finally, we include a constant α_0, a categorical variable PC_i to capture the fixed effects for each postal code area, an error term ε for area *i* in year *t*, and time dummies γ_t.

This sophisticated approach accounts for time-invariant differences between the treatment and control group. Furthermore, the year dummies eliminate any time-varying factors that affect both groups in the same way (Khandker, Koolwal, & Samad, 2009). Still, a bias in the standard error estimation can occur whenever there exists an unobserved codependence – both cross-sectional and/or autocorrelative – between one or more postal code areas within one of the groups (Moulton, 1990). To prevent such a problem to occur, we cluster our standard errors per year on postal code level (e.g. Bertrand, Duflo, & Mullainathan, 2004).

4 – Analysis

After an exploration of the data and a description of the conceptual framework in the previous chapter, we now continue to our empirical analysis. First, we describe the selection procedure that we apply to isolate both the treatment and control group from the dataset. Thereby we carefully examine the common trend assumption, which is the crucial mechanism that allows us to identify an accurate estimation of the causal effect. After this elaboration, we start our empirical analysis in Section 4. In this section we present the estimation results of the difference-in-difference models. Afterwards, we interpret and discuss these results in Chapter 5.

4.1 – Treatment and Control Group

Our dataset shows the proportion of connections within a postal code area that have a smart meter device installed. Measurement occurs once a year, on December 31st. Hence, we can construct a variable that measures the yearly implementation rate, simply by subtracting last years' percentage of smart meter connections – i.e. at the start of the year – from the percentage at measurement point. We use this implementation rate to distinguish the treatment group and control group.

As mentioned in the previous chapter, the yearly average implementation rate has been relatively constant at roughly five percent point in 2013 and 2014. Interestingly, in 2015 – i.e. the first year of the policy enactment – we only find a minor increase in the average implementation rate to 6.7 %, whereas in 2016 this rate almost doubles to 12.4 %. This suggests that we should aim our attention at 2016 as the year of treatment. When we thereby focus on the observations that have at least 50 % of smart meters placed within a specific postal code area, we find 266 areas that meet this criterion in 2016, compared to only 67 in 2015. We deliberately did not take into consideration the 317 and 396 areas with an implementation rate of at least 50 % for 2017 and 2018, respectively, because of the absence of observations for 2019. As a result, the period after the treatment would have become too short or even absent for those years. Hence, we take these 266 areas from the year 2016 as our starting point in the selection of the treatment group. As we noticed, it is important that the treatment and control group should be similar to each other. This also applies to the pre-treatment average proportion of smart meters within both groups. Hence, we apply a second criterion that there should be no more than 20 % of smart meter connections within a postal code area at the beginning of 2016. The resulting group of 200 postal code areas shows a slightly increasing trend in the average smart meter percentage before 2016 that is just below the average depicted in Figure 3.

To consider an implementation rate of 50 % to be appropriate for treatment is somewhat arbitrary, but the justification is that it provides the treatment group with a very important characteristic. As a result, the average percentage of smart meter connections within these 200 areas raises to 78.1 % after the treatment, and subsequently increases to more than 80 % during 2017 and 2018. This is exactly what the policy aims to accomplish: to equip at least 80 % of all small-scale connections with a smart meter device by the end of 2020 (Van Gerwen et al., 2010). Hence, it helps us to find the true effectiveness of the policy in its aim to reduce small-scale energy consumption by 3.2 %. Moreover, if our hypothesis is true, an implementation rate of 50 % should be large enough to discover at least some reduction in the average electricity consumption.

Regarding the selection of the control group, we can identify two important criteria. We already mentioned the first one: the pre-treatment trend in smart meter connections should be comparable between treatment and control group. The second criterion is that the treatment does not affect the smart meter percentage in the control group, meaning that the trend in this group behaves in a similar fashion post-treatment. We obtain the control group by selecting the 242 postal code areas in which the proportion of smart meters does not exceed 15 % in 2016, 20 % in 2017 and 25 % in 2018. These benchmarks allow the smart meter percentages in the control group to further develop in line with the average pre-treatment trend (Figure 4).

Before we assess the similarity of the two groups, we first remove four postal code areas that show measurement inconsistencies that persist over time. For example, an area in Amsterdam with only five residencies reported over the entire period, but where the registered number of households in that area equaled 95. In addition, we also remove 38 areas that do not provide evidence on key variables – e.g. income, year of construction, various population characteristics – over time. These observations are mostly on industrial areas where the number of residents is so low, that the CBS does not publish accurate information for that area due to privacy concerns. What further defines these areas is a very high level of energy consumption – i.e. 8,495 kWh on average – and a large number of firms per inhabitant. Appendix D shows the distribution in electricity consumption before and after adapting the sample. The result is a strongly balanced sample with a treatment group that contains 191 postal code areas, compared to 213 areas in the control group.

Figure 5 depicts the resulting average percentage of smart meters for both the treatment and control group. Clearly visible is the constant development in the percentage of smart meter connections for the control group and the instantaneous increase for the treatment group after 2015. In addition, the gradual

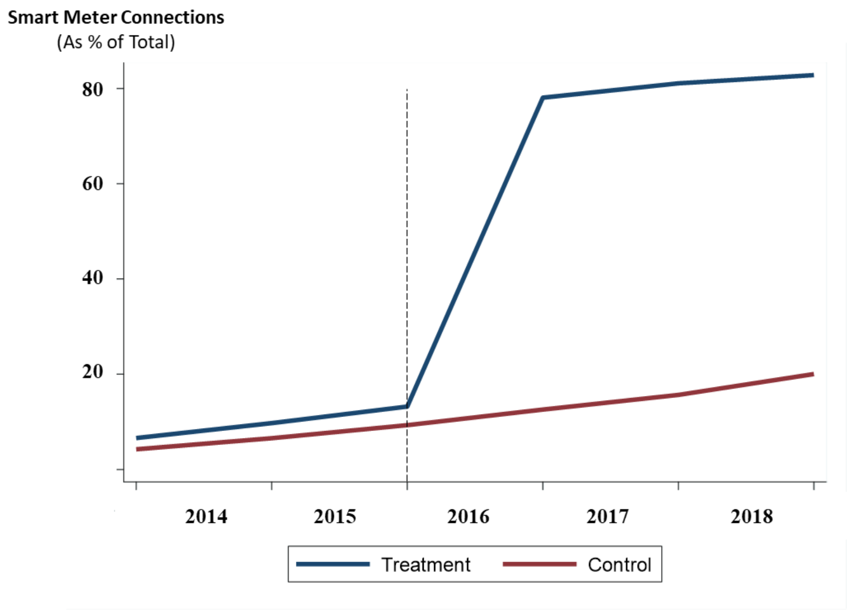

FIGURE 5 - Development in the Percentage of Average Smart Meter Connections within the Sample

development in the implementation of smart meters in the treatment group seems to persist again after the treatment took place. The average implementation then equals 2.4 percent point per year in this group, which hardly differs from the yearly average implementation rate of 3.2 percent point in the control group. What may explain this development is the fact that, under normal conditions, all newly constructed buildings receive an active smart meter connection.

4.2 – Common Trend Assumption

In the previous chapter we already explained how the difference-in-difference methodology accounts for possible differences between the treatment and control group. However, there is one crucial requirement that we further need to examine in order to make this framework successful: the common trend assumption (e.g. Angrist & Pischke, 2014) – also referred to as parallel trends assumption (e.g. (Khandker, Koolwal, & Samad, 2009). The idea of this assumption is that, in absence of treatment, the average electricity use in the treatment group would have evolved in the same way as it actually evolved in the control group. There are two reasons why we can justify this assumption. First, the fact that the two groups show similar average values on all observable characteristics – i.e. the four categories we described in Section 3.3. Table 1 displays figures on the most important control variables over the entire period. In this comparison between the two groups, we discover that there are no significant differences present over the years between the treatment and control group. Still, it appears from the table that an area within the treatment group contains a higher number of inhabitants – and hence more connections and residencies – on average. However, this difference is far from significant as well. Moreover, the population density score appears to be almost identical between the two groups, which therefore means that the space area is larger in the treatment group on average. More important is the fact that the number of connections within an area – regardless which group it belongs to – hardly differs from the number of residencies on average. More specifically, areas in both the treatment and control group consist for more than 95 % of residential connections, which illustrates that our sample primarily contains areas that are residential in nature. Furthermore, this correspondence resolves the limitation that we cannot identify the different types of small-scale users, since both groups appear to be identical in the proportion of non-residential small-scale users. Ultimately, the treatment status did not affect any of the control variables from 2016 onwards, which also strengthens the validity of the assumption (Khandker, Koolwal, & Samad, 2009).

The second reason is an identical trend in average electricity use that we find in both groups before the treatment. Figure 6 clearly shows that, before 2016, the average consumption hardly differs between both groups and develop in a similar trend. Combined, on average, this consumption is 102 kWh lower than the consumption average in the total dataset (Figure 3). Hence, based on the observables and the similarity in trends, we can infer that both groups are the same, apart from their treatment status. Therefore, it is reasonable to assume that the electricity consumption in both groups reacted equivalently to – either visible or invisible – changes after the treatment as well. In other words, we can use the trend in the control group as the counterfactual for the treatment group.

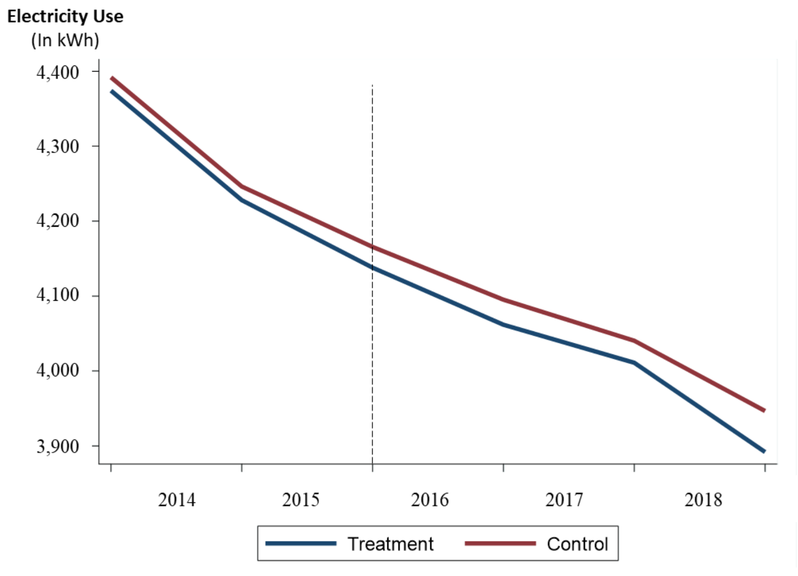

FIGURE 6 - Development in the Average Use of Electricity across Treatment and Control Group

A very convenient implication follows from this assumption. In the previous chapter, we discussed how we included various factors that serve as an approximation to control for characteristics that are hard to measure. For example, the average energy label values to approximate the improvements in energy performance of buildings. Thereby our conclusion was that we cannot control correctly for particular variables, especially the factors on individual household level related to lifestyle or the type of energy contracts. However, these concerns become irrelevant if the treatment and control group are the same on average, because it implies that not only the quantifiable variables are similar – as we verified in Table 2 – but also the unobservable characteristics. Hence, controlling for these characteristics in our analysis is not necessarily useful from the perspective of the estimator (e.g. Angrist & Pischke, 2014; Lechner, 2010). On the other hand, not including such factors induces two minor disadvantages to our analysis. First, the missing factors could have helped us to detect possible heterogeneity in the causal effect estimation. For instance, whether different types of energy contracts influence the effectiveness of the indicative overview of user cost. Second, it makes the clustered standard error estimation less precise (Angrist & Pischke, 2014).

4.3 – Credibility of the Common Trend Assumption

Table 1 shows that the treatment appears to have no effect on any of the variables that characterize our sample. This provides a solid base for the credibility of the common trend assumption. However, the Table also raises a point of concern: the percentage of connections with district heating shows a persistent dissimilarity over time. In the treatment group, 19.4 % of the connections use district heating on average, compared to only 4.7 % in the control group. However, both groups show large variation in the distribution of the average percentages of district heating across postal code areas. As a result, this difference in the average percentage is not significant, but it may still explain why the average electricity consumption is somewhat lower in the treatment group. In addition, when we compare the distribution in the percentage of district heating over time, it is easy to observe that the standard deviations hardly change over time. Moreover, the trend in the average percentage of district heating connections appears to be identical in the two groups. That is to say, these percentages are very stable over the years, apart from one sudden increase in 2015 which occurred in both groups. Since this particular trend does not differ between the groups, including them into our regression analysis prevents a possible violation of the common trend (e.g. Angrist & Pischke, 2014).

A second point of concern that may lead to a violation of the common trend assumption is the percentage of users who switched off their smart meter. Since small-scale users in the control group did not – yet – receive an offer of smart meter placement, this proportion of users is reasonably larger in the treatment group. In this regard, the treatment group might differ from the control group. Unfortunately, the percentage of smart meter connections that have been switched off is unobservable for each postal code area. As a result, we cannot control for it, but the question is whether this can affect our estimation. In 2016, only 2 % of the installed smart meters – i.e. 0.007 % of all connections in our dataset – have been switched off. By the end of our time period, this percentage barely increased to 3 % (RVO, 2018), which is so low that it can only affect our analysis if and only if a disproportionate number of users in treatment group areas switched off their smart meter. However, as Figure 7 shows, the spread in treated

TABLE 1

Means and Standard Deviations on the Variables across Treatment[a] and Control[b] Group over Time.

	2013		2014		2015		2016		2017		2018	
	Treatment	Control	Treatment	Control	Treatment	Control	Treatment	Control	Treatment	Control	Treatment	Control
Electricity Use (kWh)	4,374.13	4,391.96	4,228.03	4,246.24	4,137.91	4,165.59	4,061.77	4,095.20	4,011.06	4,040.46	3,891.89	3,946.63
	(1,204.53)	(1,080.56)	(1,152.27)	(1,037.13)	(1,129.26)	(1,025.84)	(1,097.98)	(1,002.71)	(1,094.00)	(1,002.68)	(1,072.90)	(977.62)
- Avg. Delivery (%)	97.6%	97.8%	96.9%	97.0%	95.9%	96.1%	94.8%	95.3%	93.5%	94.3%	91.4%	91.9%
	(2.71)	(1.85)	(3.04)	(2.18)	(3.38)	(2.66)	(3.81)	(3.00)	(4.35)	(3.56)	(5.22)	(4.47)
Number of Connections	2,428	1,715	2,435	1,718	2,447	1,721	2,468	1,730	2,510	1,761	2,503	1,745
	(1,995)	(1,861)	(1,995)	(1,862)	(2,007)	(1,868)	(2,018)	(1,879)	(2,057)	(1,919)	(2,037)	(1,895)
- % Smart meters	6.6%	4.2%	9.7%	6.6%	13.2%	9.3%	78.1%	12.6%	81.1%	15.7%	82.8%	20.0%
	(2.66)	(1.47)	(3.18)	(1.83)	(3.65)	(1.81)	(7.20)	(2.05)	(6.23)	(2.57)	(5.48)	(3.17)
- % Heavy Connections	0.4%	0.4%	0.4%	0.4%	0.4%	0.4%	0.4%	0.4%	0.4%	0.4%	0.4%	0.4%
	(0.22)	(0.20)	(0.22)	(0.20)	(0.22)	(0.20)	(0.20)	(0.20)	(0.22)	(0.20)	(0.22)	(0.20)
- % District Heating	15.9%	3.6%	16.1%	3.5%	20.9%	5.2%	21.2%	5.4%	21.2%	5.3%	21.3%	5.3%
	(33.74)	(14.47)	(33.51)	(14.35)	(35.57)	(18.49)	(35.58)	(18.44)	(35.08)	(18.25)	(35.23)	(18.33)
Gas Consumption (m^3) *	1,791.84	1,792.70	1,736.46	1,731.09	1,697.24	1,624.79	1,680.73	1,674.45	1,655.02	1,677.81	1,650.10	1,671.52
	(629.11)	(563.15)	(633.19)	(535.49)	(634.15)	(517.00)	(622.91)	(513.35)	(596.99)	(513.35)	(639.64)	(511.36)
Number of Inhabitants	5,449	3,708	5,473	3,709	5,387	3,638	5,412	3,640	5,442	3,648	5,477	3,658
	(4,561)	(4,212)	(4,613)	(4,205)	(4,392)	(4,131)	(4,408)	(4,134)	(4,430)	(4,147)	(4,455)	(4,161)
- Avg. Age	39	40	39	40	40	41	40	41	40	41	40	41
	(3.28)	(3.91)	(3.28)	(3.83)	(3.26)	(3.71)	(3.22)	(3.63)	(3.17)	(3.54)	(3.18)	(3.48)
- % Male	49.6%	49.4%	49.7%	49.4%	49.8%	49.5%	49.8%	49.4%	49.8%	49.5%	49.8%	49.5%
	(1.61)	(1.57)	(1.60)	(1.56)	(1.56)	(1.55)	(1.54)	(1.51)	(1.46)	(1.48)	(1.49)	(1.51)
- % Dutch	76.2%	78.4%	75.8%	78.2%	76.3%	78.8%	75.6%	78.5%	75.0%	77.8%	74.6%	77.3%
	(14.14)	(16.51)	(14.22)	(16.47)	(14.99)	(17.32)	(15.56)	(17.47)	(15.62)	(17.69)	(15.69)	(17.87)
- % Active	57.8%	59.4%	57.7%	58.6%	57.6%	58.4%	58.0%	58.9%	58.5%	59.2%	59.5%	60.2%
	(8.31)	(7.62)	(7.06)	(7.27)	(6.84)	(7.12)	(6.70)	(6.95)	(6.53)	(6.73)	(6.43)	(6.53)

Mean values as of December 31st; yearly standard deviations in parentheses and denoted in percent points for the variables measured in percentages. [a] n = 191; [b] n = 213; * 29 postal code areas in the treatment group do not have any gas connections, compared to 4 areas in the control group.

TABLE 1 (CONT.)

Means and Standard Deviations on the Variables across Treatment and Control Group over Time.

	2013		2014		2015		2016		2017		2018	
	Treatment	Control	Treatment	Control	Treatment	Control	Treatment	Control	Treatment	Control	Treatment	Control
Number of Households	2,436	1,635	2,443	1,635	2,422	1,618	2,376	1,623	2,394	1,632	2,416	1,637
	(2,128)	(1,852)	(2,143)	(1,843)	(2,085)	(1,826)	(2,100)	(1,828)	(2,112)	(1,840)	(2,144)	(1,847)
- Avg. Household Size	2.22	2.24	2.22	2.24	2.21	2.23	2.21	2.22	2.20	2.21	2.19	2.20
	(0.32)	(0.32)	(0.31)	(0.32)	(0.31)	(0.32)	(0.30)	(0.31)	(0.30)	(0.31)	(0.31)	(0.31)
- Avg. Income (€ 1,000)	21.979	22.744	22.643	23.379	23.607	24.614	24.125	25.307	25.084	26.231	24.782	26.177
	(3.120)	(5.330)	(3.349)	(5.825)	(3.980)	(7.157)	(3.960)	(7.242)	(4.823)	(7.830)	(4.699)	(7.703)
Number of Residencies	2,361	1,645	2,389	1,660	2,362	1,636	2,376	1,642	2,394	1,646	2,416	1,652
	(2,009)	(1,822)	(2,038)	(1,834)	(1,979)	(1,815)	(1,988)	(1,822)	(1,997)	(1,827)	(2,017)	(1,834)
- Year of Construction	1975	1972	1975	1972	1975	1972	1976	1972	1976	1972	1976	1972
	(14.29)	(14.85)	(14.32)	(14.86)	(14.34)	(14.88)	(14.25)	(14.86)	(14.14)	(14.99)	(14.05)	(14.89)
- % Rental	44.2%	41.8%	42.8%	40.6%	42.1%	40.6%	42.3%	40.4%	42.4%	40.6%	42.3%	40.4%
	(17.32)	(14.57)	(17.67)	(15.04)	(17.20)	(14.15)	(17.33)	(14.34)	(17.18)	(14.36)	(17.15)	(14.33)
- % MDU's	33.5%	34.5%	33.8%	35.0%	33.7%	35.8%	33.7%	35.8%	33.9%	36.2%	33.9%	35.6%
	(24.07)	(27.62)	(23.98)	(27.52)	(23.60)	(26.82)	(23.44)	(26.83)	(23.41)	(27.06)	(23.26)	(26.76)
- Avg. Living Space	116.62	113.14	115.38	112.74	115.31	112.53	115.37	113.75	115.36	115.82	115.59	116.83
	(21.37)	(41.82)	(20.07)	(41.38)	(20.31)	(41.17)	(20.35)	(38.17)	(20.26)	(35.28)	(20.10)	(34.40)
- Avg. Energy label**	2.41	2.46	2.39	2.44	2.33	2.40	2.29	2.40	2.27	2.39	2.19	2.32
	(0.35)	(0.39)	(0.35)	(0.39)	(0.38)	(0.40)	(0.38)	(0.41)	(0.40)	(0.42)	(0.39)	(0.41)
Number of Businesses	382	278	388	283	400	293	412	303	422	309	442	325
Density* **	1,719	1,693	1,737	1,692	1,771	1,708	1,780	1,716	1,793	1,726	1,804	1,737

** Average value of valid labels, on scale 1 (i.e. best, A-label) to 4 (i.e. worst, F- or G-label); *** Measured in number of addresses per ¼ square kilometer.

postal code areas across the country makes any of such inferences unreasonable. On top of that, the difference in the percentage of smart meters that have been switched off does not harm our analysis.

We just described that both groups are not similar regarding the district heating connections, and that we have to control for this dissimilarity. Let us consider this point from a different perspective. What if the net managers prioritized postal code areas with a high proportion of district heating connections in their implementation efforts? If this argument holds, the common trend assumption remains valid. However, it may indicate that the policy has been applied selectively to such areas, which distorts the exogeneity of the policy and subsequently causes an undesired correlation between district heating and smart meter percentages. To find out whether this is the case, we first describe the development from policy enactment to the large-scale implementation of smart meters. After that, we provide several arguments that validate the exogeneity of the policy, and we elaborate more on the suggested presumption.

As we described in Chapter 2, the Dutch government enacted the implementation policy in 2012 with the aim to equip at least 80 % of all connections with a smart meter device by the end of 2020. Based on earlier calculations by Van Gerwen et al. (2010), regularly informing consumers on their energy consumption behavior could result in a 3.2 % reduction in small-scale electricity use on average. Consequently, consumers benefit from a smart meter through a lower electricity bill. But are they also responsible for purchasing and installing a smart meter device?

This would unnecessarily complicate things. Therefore, after the policy enactment in 2012, the government made the net managers responsible for the placement of the smart meters. Hence, the net managers have to purchase the devices from a producer and subsequently install them on their own initiative (Vringer & Dassen, 2016). Net managers thereby have to offer a smart meter device to all small-scale consumers for free. In return, net managers earn back this investment through the meter rent which users pay as part of the yearly electricity bill. Consumers receive this offer in the form of an official letter which users receive three months prior to the upcoming installation, hence reducing the possibility that consumers develop any kind of anticipatory behavior in their consumption of energy (e.g. Heckman, Lalonde, & Smith, 1999). Consumers can refuse the offer, but if they do so, any future installation is at their own cost. Hence, if he/she changes opinion or if the old electricity meter breaks down after the offer has been rejected, the net manager is allowed to charge the full cost of € 72,60 to the consumer (ACM, 2019).

At first, the net managers started their implementation efforts according to some sort of trial. Newly constructed buildings received a smart meter, as well as existing buildings that underwent major renovations or in case the old electricity meters had to be replaced (De Lange, 2018). After these experiences, the government decided to further expand the implementation program. Hence, from 2015 onwards, the efforts have been extended to all small-scale users across the country. Thus, by contracting out the duty to install the smart meters, the government did not design a specific order or protocol which areas or buildings had to be treated first. This already provides a weak indication why we can consider the implementation policy as exogenous. To validate their independent responsibilities in their implementation efforts, Enexis, Liander and Stedin provide overviews of the treatment status per four-digit postal code area and the year and quarter in which treatment occurred (Enexis, 2020; Netbeheer Nederland, 2020). When we sort the observations by year and by the quarterly implementations, we discover a random spread in the treated postal codes across the different provinces in the Netherlands within each quarter. This is also logical from a staffing perspective, since it is not very efficient to let all the installation technicians work in only one district at the time. Figure 7 shows a map of the Netherlands that categorizes all 4,069 postal code areas by their percentage of smart meter connections as from the 1st of April 2020. Even a year after our period ended, the random distribution in the clusters of treated- and non-treated areas is clearly visible. Thereby we do not find any evidence that net managers have been prioritizing densely populated areas over rural areas, since many areas in Gouda, The Hague, Arnhem, Leeuwarden and several other large cities still await a smart meter implementation on large scale. When we further specify the criteria to the treatment group in our sample, we obtain the same conclusion. In Figure 7, we visualize the random assignment of treatment by means of blue dots to indicate the exact location of the 191 treatment areas. Since Enexis, Liander and Stedin do not serve the Zeeland province, we do not find any observations there.

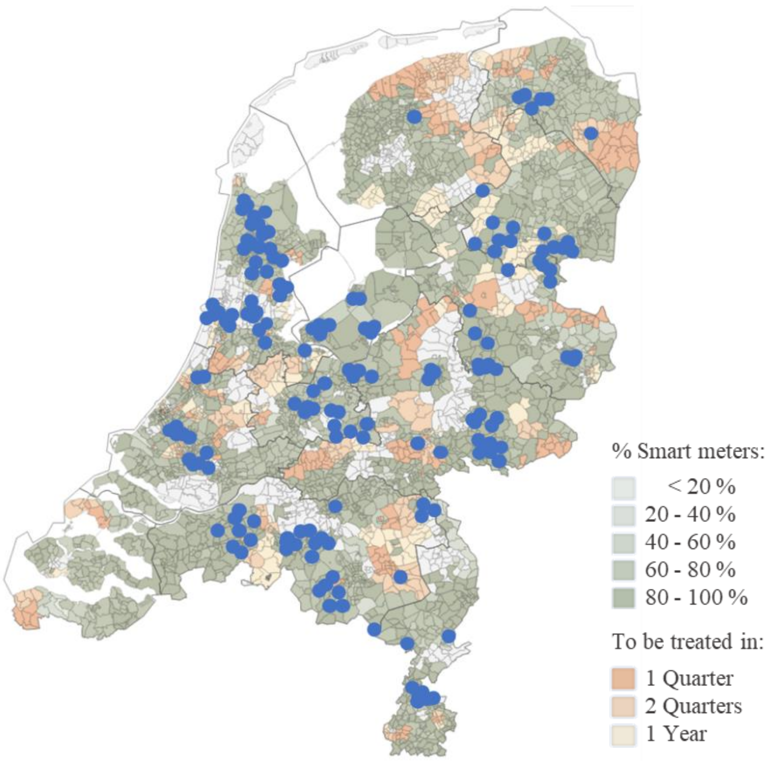

FIGURE 7 - Spread in Smart Meter Connections across the Netherlands

However, if the treatment has been assigned randomly to postal code areas, why then is the percentage of district heating connections persistently higher in treated areas? To understand the reason behind this phenomenon we first note that, in the Netherlands, the use of a district heating infrastructure is not only limited to cities or city agglomerates. In 2016, a total of 356 postal code areas across the entire Netherlands utilized this source of energy, of which 81 areas comprised small cities and even villages. However, when we consider areas in which more than half of all connections make use of city heating, it becomes clear that 72 % of these areas are located in central provinces Utrecht and Flevoland, and western provinces Noord- and Zuid-Holland. Hence, city heating is strongly represented in the central-western region relative to the other parts of the Netherlands. This fact may cause the problem we want to understand, since Figure 7 clearly shows that there are several – clusters of – treatment areas located in those four provinces.

As we showed in Figure 1, both Liander and Stedin serve this central-western region. And exactly these two net managers account for the imbalance between the groups. After we categorize treatment and control group by net manager, we find that the average district heating percentage is twenty percent point higher in the treated areas served by these two net managers. Straightforwardly, because the treatment group contains 22 areas located in the central-western provinces and in which more than half of the connections utilize district heating. Since our entire sample contains only 53 areas – 42 treatment against 11 control areas – where the district heating percentage is larger than zero, it is easy to see that observations from this part of the country strongly influence the difference between the two groups. Conversely, the ten treatment and control areas served by Enexis show a slightly imbalanced district heating percentage in favor of the control group – 6.7 % versus 10.5 %, respectively.

TABLE 2
Number of Treated District Heating Areas per Year

	2015	2016	2017	2018	Untreated [1]
Enexis	2	2	2	0	4
Liander	1	24	1	1	4
Stedin	7	4	2	0	2

To see whether Stedin and Liander prioritized the implementation of smart meters in areas where the majority of connections includes a district heating connection, we compare the number of such areas that received treatment over time. In 2016, Stedin only treated four of such areas, which equals just 14 % of their total treatments of that year. Clearly, these four treatments do not stand out from the number of district heating areas treated by Stedin in other years. Table 2 compares the number of areas in our entire dataset in which the percentage of district heating connections exceeds 50 % and in which a large-

[1] Areas in which the majority of addresses have a district heating connection, but where no major hike in smart meter implementation occurred – i.e. 3.8 % yearly average implementation, and less than 50 % smart meters by the end of 2018.

scale implementation of smart meters – i.e. more than fifty percent point – occurred within a specific year over our time period. Hence, based on observations in other years, we can confidently claim that Stedin did not prioritize areas where district heating is strongly present. For Liander, however, we cannot make such an inference. Obviously, the number of 24 district heating areas in 2016 shows how Liander prioritized areas with a large share of district heating connections over areas in which gas connections are more common. Unfortunately, our treatment group includes twenty of those areas, most of them clustered around the cities of Alkmaar, Almere and Purmerend. What disturbs the balance between our treatment and control group even more is the fact that, in these areas served by Liander, the average proportion of district heating connections equals 95.7 %. Hence, not all three net managers have been randomly placing the smart meters on a large scale, which partly violates the policy exogeneity. In the next section, we show how we solve this correlation issue.

4.4 – Regression Analysis

In Figure 6 we can already observe that the treatment hardly affected the average energy consumption in the treatment group. After 2016, both lines seem to continue their identical trends, which may indicate that the smart meter implementation did not have an effect at all. Table 3 summarizes the levels and changes in average consumption of electricity per connection within both the treatment and control group. The data in the first column represents the average consumption per connection for areas in the control group, and the data in the second column represent those for the treated areas. In addition, before refers to average energy consumption over the years prior to the implementation, and after refers to the consumption over the years 2017 and 2018. The third row of the table presents the differences between the treatment and control group, both before and after the treatment. This allows us to make a basic estimation of the treatment effect, simply by subtracting the average decline in consumption in the control group from the average decline in the treatment group (e.g. Angrist & Pischke, 2014; Card & Krueger, 1994). The relative decrease in average energy consumption is thus equal to 20.79 kWh, which is not significant.

TABLE 3

Average Energy Use per Connection, Before and After the Smart Meter Implementation

	(I) **Control**	(II) **Treatment**	(III) **Difference**
1. Before	4267,80	4246,39	-21,41
2. After	3993,76	3951,56	-42,20
3. Change in avg. consumption	-274,04	-294,83	-20,79

Moreover, is this effect entirely attributable to the large-scale implementation of smart meters? To see whether there is an actual decrease, we now continue to empirically test the effect of the implementation on the average consumption of electricity. Table 4 presents the results of the difference-in-difference analysis. In order to accurately test the hypothesis, we first run a baseline specification with only the treatment dummy, together with the year- and area fixed effects. The resulting estimate is slightly lower compared to the basic exploration of the first estimation model in Table 3 and shows a weak significance. This inconsistency between the two estimates emerges because we cluster the observations on area- instead of group level. Moreover, it may also expose a bias in our simple comparison of averages that causes an underestimation of the true effect. As explained, we expect that the dissimilarity in district heating connections between the treatment and control group causes this bias. To discover the magnitude of the correlation, we add only the district heating percentage to the basic exploration. The resulting estimation of the treatment effect in the second model decreases by almost fifteen percent, which is a clear indication of a correlation between the two variables.

Therefore, our more extensive models all include the district heating percentage on postal code level. In addition, we add the average living space and the eight categories for the year of construction as additional control variables. Although these variables do not significantly differ between the groups on average, we include the former because of the large difference in standard deviations, and the construction categories based on the findings of Aksoezen et al. (2015). First, we run a regression with the control variables in which we exclude the treatment dummy (Column III). Subsequently, we include the dummy as well, and compare the outcomes of the estimators between the second, third and fourth model. Again, the correlation between the district heating estimator and the treatment effect is clearly visible. After leaving out the treatment dummy in the third model, the district heating estimator decreases considerably as a result. More important, however, is the fact that the treatment estimation is almost equal in the second and fourth model. Another interesting finding is that the estimators of the control variables displayed in Columns III and IV hardly differ. More important, however, is the fact that the treatment estimation is almost equal in the second and fourth model. These two findings provide clear evidence that there is no correlation between the estimated treatment effect and the other control variables.

What remains is the correlation between district heating and the smart meter implementation. To solve this correlation bias, we add an interaction term between the treatment variable and the district heating percentage (e.g. Bun & Harrison, 2019). In this way we allow for pre-treatment differences in district heating across particular areas within the treatment group. Column V in Table 4 shows the estimation results of this interaction model. The estimator for district heating is now almost the same as the estimate that results from the control regression (Column III), which indicates that the interaction adequately resolves the correlation. Moreover, compared to the third model, the interaction term attributes an additional increase to the treatment effect.

TABLE 4
Estimation Results of the various Difference-in-Difference Models

	(I)	(II)	(III)	(IV)	(V)	(VI)
Treatment	-29.8*	-34.204**		-35.609**	-39.877**	-41.646**
	(17.438)	(17.407)		(17.091)	(19.484)	(19.250)
District Heating		187.450**	155.067**	170.498***	159.350**	134.591***
		(64.366)	(61.044)	(62.677)	(65.961)	(66.383)
Avg. Living Space						
- SDU's			-0.978**	-1.053**	-1.046**	-0.964**
			(0.420)	(0.442)	(0.442)	(0.432)
- MDU's			1.057**	1.053*	1.037*	0.942*
			(0.516)	(0.558)	(0.559)	(0.541)
Constant	3,626.3***	3631.857***	5,025.7***	5,002.935***	4,994.466***	4,516.641***
	(9.187)	(9.428)	(546.520)	(581.956)	(568.602)	(59.167)
Net Manager	No	No	Yes	Yes	Yes	Yes
Construction categories	No	No	Yes	Yes	Yes	Yes
Time Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
Area Fixed Effects	Yes	Yes	Yes	Yes	Yes	Yes
Interaction term	No	No	No	No	Yes	Yes
Observations	2,424	2,424	2,424	2,424	2,424	2,304
R^2	0.991	0.991	0.991	0.991	0.991	0.991

Notes: * denotes significance at the 10 % level; ** significant at the 5 % level; *** significant at the 1 % level. The denoted estimators are measured in terms of kWh electricity use. Robust standard errors clustered by postal code area (404 clusters) and reported in parentheses. Net Manager is a categorical variable that indicates which manager serves the postal code area.

Two findings enhance the credibility of this final regression. First, we run the fourth model again, but thereby excluding the twenty prioritized areas in which the district heating percentage equals 95.7% on average. Column VI depicts the result of this supplementary regression. Clearly visible is the same shift in effects – which we already observed in Column V – between the treatment dummy and the district heating variable. Second, adding any other control variable in the fifth model hardly affects the treatment effect estimation. This is consistent with our expectations based on Table 1 which shows that there are no major differences on average between the treatment and control group. Hence, we can confidently state that there are no other correlations that could bias the results. The results of this final regression show that we cannot reject our main hypothesis that the smart meter implementation on large scale reduces the yearly average electricity consumption (P = 0.041). More specific, if a net manager equips the majority of all small-scale connections within an area with a smart meter, the yearly electricity consumption for a small-scale user in that area decreases by 39.9 kWh on average.

Note that the explanatory power for all models is extremely high ($R^2 = 0.991$), because we added the fixed effect dummies for each individual postal code. In this way, our specification captures all unobservable time-invariant effects for each specific area. This implicitly means that we run a specific regression for each postal code area instead of estimating one, generic regression for all observations in one of the two groups. Clearly, this specification is a better fit to the data, thereby leading to much smaller residuals (Angrist & Pischke, 2014). If we include a group dummy instead of the area fixed effects, the R-squared value equals a more moderate 0.301 in the fourth model.

4.5 – Robustness Check

In this final section of our analysis, we conduct a robustness check to examine the validity of our estimation (Khandker, Koolwal, & Samad, 2009). We do so by running a panel regression on the same observations, but now with the smart meter percentage as the independent variable. Hence, we now estimate the effect of an increase in the proportion of smart meter connections instead of the effect of a binary treatment variable. This robustness check validates our main finding if we find a negative association between the smart meter percentage and the average electricity consumption. We estimate the relationship between electricity consumption and the smart meter percentage for the entire sample only – instead of two regressions for the treatment and control group separately. We do so, since the lack of district heating areas in the control group causes a strong bias in the district heating estimate and the interaction term. Appendix E displays the results of the fixed effects regression. As we already mentioned in the previous chapter, we are not able to include several important variables in this regression, which implies that an absolute interpretation is not appropriate. However, this model still provides valuable information, since it gives a weak indication of a small and negative correlation between the proportion of smart meters and the average electricity consumption.

5 – Discussion

This chapter presents an elaborate discussion on the empirical results we presented in Section 4.4. We thereby start with an interpretation of the results in the first section. After that, we discuss the implications that follow from our findings, and derive a recommendation for the Dutch government. We end this chapter by considering the limitations of our research and suggesting some possibilities for future research.

5.1 – Interpretation of the Results

We started our empirical analysis by asking what the effect of a large-scale implementation of smart meters on the average electricity consumption has been. Thereby we mentioned prior research that showed how a smart meter device could help to reduce the energy consumption of households. The mechanism that effectuates this reduction, is providing energy consumption feedback to the users. To see whether this mechanism also holds within the context of a large-scale implementation of smart meters, we employed a dataset containing cross-sectional observations on almost all four-digit postal code areas within the Netherlands over a time period from 2013 until 2018. Based on the information this dataset provides, we discovered that in areas where such an implementation took place, small-scale users in these areas subsequently reduced their yearly electricity consumption with 39.9 kWh on average. It is important to notice that we have to interpret this finding carefully, because the dataset reports the measure of electricity consumption as the average per-connection consumption for all connections within a postal code area. Hence, we must not presume that the implementation of smart meters reduces the electricity consumption for each individual user, simply because not all connections within a treated postal code area received a smart meter. Consequently, the true effect for an individual user who received a smart meter might be even stronger.

Why is this finding so important then? To have a correct understanding, we have to bring in mind the intention of the policy to equip at least 80 % of all connections with a smart meter. The government estimated that, if this intention has been accomplished, a reduction in small-scale electricity consumption by 3.2 % could be realizable. From our selected sample emerges that the average proportion of smart meter connections in the treated areas increased to more than 80 % on average as a result of the implementation efforts. The effect we found is thus useful in evaluating whether the implementation efforts had the desired effect. Straightforwardly, we can compare the measured reduction of 39.9 kWh with the pre-treatment average electricity consumption of 4,137.9 kWh in the treatment group, which implies a treatment effect of -0.96 %. Thus, to summarize, we cannot conclude that placing a smart meter device – and consequently informing the user – reduces the electricity consumption of this user by 39.9 kWh. What we can conclude, however, is that the implementation policy helped to reduce the average electricity consumption of an average small-scale user by almost one percent, which is a much weaker effect than expected.

5.2 – Implications and Recommendations

We can derive some meaningful academic and practical implications from this study. First, the fact that providing bi-monthly consumption feedback leads users to reduce their energy consumption serves as a validation of prior research. However, the measured reduction of 0.96 % is much weaker than the estimated effects that emerge from previous researches. As discussed in Section 2.3, these studies have mostly been based on experimental settings and have all been carried out on a smaller scale compared to our study. As a result, the true effect of a large-scale implementation is much weaker than expected, since the expectations of Van Gerwen et al. (2010) and the Dutch government have been based on estimated effects in experimental settings.

Three reasons explain why the measured effect does not match these expectations. First of all, people are known to behave differently in experimental settings – called the Hawthorne effect (Franke & Kaul, 1978; Mayo, 2004). In other words, people are more likely to behave according to the expectations when they are involved in an experiment, whereas this willingness disappears in the context of the entire society. Second, and related to the first reason, not all users may be interested in the received feedback. For example, users do not take the time to read the indicative overview of user cost, do not understand what the overview says, or are simply not willing to change their consumption behavior for many reasons. The third reason why the measured effect does not match the estimations from experimental settings is because the latter measurements have been based on a binary condition: treated users receive a smart meter, whereas untreated users do not have a smart meter. However, the implementation efforts in our study are not binary in nature, since our unit of observations are the postal code areas in which the smart meter measure is a percentage of the total connections within that area. We thereby selected the areas with an increase of at least fifty percent point in the percentage of smart meters, which obviously lessens the strength of the treatment compared to the binary smart meter status in experimental settings. Moreover, the smart meter measure in the control areas is not equal to zero either. As a result, many users do not comply – unintentionally – to the treatment status of the area in which they live or work (Angrist, Imbens, & Rubin, 1996).

Hence, in practice, it turns out that the effect of equipping small-scale users with smart meters is lower when implemented on a large scale. We can therefore derive a meaningful practical implication as well. Since the energy savings attainable from this type of information provision is only 0.96 % on average, the Dutch economic assessment of the smart meter implementation has been too optimistic. Other Member States of the European Union should therefore be cautious not to overestimate the benefits of a countrywide installation of smart meters. Furthermore, when thinking one step ahead, we can also derive a recommendation for the Dutch government in particular. As we described in the second chapter, the government justified the smart meter implementation policy primarily because of the expected reduction in energy consumption that smart meters facilitate (Van Gerwen et al., 2010; De Lange, 2018; RVO, 2018). However, by emphasizing the importance of reducing the consumption of electricity, the implementation policy reveals a major discrepancy in the government's intentions to actually reduce the total electricity use in the Netherlands. The cause of this discrepancy is the fact that small-scale electricity use comprises only 20 % of the total electricity use in the Netherlands (Hieminga, 2013). Now, illogically, net managers have spent an estimated € 3.3 billion (Kamp, 2014) – which will mostly be paid by the consumers in the long run – to realize a very optimistic reduction of 3.2 % that only applies to just 20 % of the total electricity use. Consequently, if reducing electricity use is such an important objective for the Dutch government, why did it not implement reduction measures that apply to the large-scale business users responsible for the other 80 % of the total electricity use? The Dutch government should therefore put a stronger focus on the large-scale business users in their efforts to effectively reduce the consumption of electricity in the future.

5.3 – Limitations and Future Research

In Chapter 2 and Appendix A, we showed that the indicative overview of user cost has been standardized according to the Dutch Smart Meter Requirements. Consequently, in our analysis, we more or less assumed that the feedback provided to small-scale electricity users entailed homogenous information. However, there may still be heterogeneity in the type and frequency of the provided feedback for individual users. We already mentioned the fact that some users switched off their smart meter device. In addition, we did not take into account the users who received direct feedback from an In-Home Display. Although the percentage of all small-scale users with an In-Home Display was only equal to 4 % in 2018 (RVO, 2018), it still resulted in different types of feedback on electricity use that we could not control for. This heterogeneity in feedback is the first limitation of our analysis, since it could have influenced the measured effect on the average electricity use. The second limitation is that we have not been able to distinguish between the different types of users and their individual behavioral characteristics. As a result, we could only estimate a general treatment effect, without discovering any heterogeneity in the effect of indirect feedback among individual users.

Based on these limitations, we derive some suggestions for future research. First of all, it can be useful to gather additional data regarding the sources of feedback within each postal code area, and subsequently investigate the effectiveness of indirect feedback in the presence of direct feedback. In addition, it may be interesting to study other sources of heterogeneity as well. For example, to investigate whether the estimated effects vary across different cities or provinces. Or alternatively, whether different types of small-scale users respond differently to feedback on electricity use, and what behavioral characteristics cause these different reactions. Investigating these possible relationships provides better insight in the effectiveness of feedback on electricity use, which may help to improve the distribution of feedback in the future. Second, adding observations on the variables over 2019 can further enhance our analysis, since it offers the possibility to select treatment areas over different years. However, this in turn complicates the control group selection; a more sophisticated method called Matching (e.g. Khandker, S.R.; Koolwal, G.B.; Samad, 2009) facilitates this selection procedure. A final suggestion is to investigate whether the described effect of a large-scale implementation persists in the long run as well.

6 – Conclusion

Both in our study and in earlier research, it appears that the provision of feedback can significantly contribute to reductions in electricity use. In theory, regularly confronting consumers with their electricity use creates better awareness, which subsequently activates consumers to reduce their electricity use. In addition, previous literature described several factors that determine the effectiveness of this feedback mechanism, such as the type of feedback and the differences in feedback strategies.

In this master thesis, we studied the effectiveness of providing indirect feedback to small-scale users in the Netherlands. Each small-scale user who has been equipped with a smart meter connection received this feedback through a bi-monthly indicative overview of user cost. In order to utilize this feedback mechanism, the Dutch government enacted that at least 80 % of all small-scale electricity users should be equipped with a smart meter connection by the end of 2020. The government's primary justification for this implementation policy has been an estimated 3.2 % decline in the aggregate small-scale electricity use. In our analysis, we empirically verified the accuracy of this estimation through using a dataset that contains observations on the average electricity use per connection within a postal code area. From this dataset, we selected a total of 404 areas that showed very similar characteristics over time, apart from their percentage of smart meter connections. In 191 postal code areas, a large-scale implementation of smart meters took place in 2016, hence leading to an 80 % smart meter connection average within these areas. In the other 213 areas, no such implementation occurred over time. After selecting these two groups, we used several difference-in-difference estimations to calculate the effect of the large-scale implementation of smart meters in terms of the average electricity use. Interestingly, the measured reduction equaled 0.96 %, which is much lower than the expected reduction of 3.2 %. Hence, the policy has not been very effective in reducing the small-scale use of electricity in the Netherlands.

References

Abrahamse, W., Steg, L., Vlek, C., & Rothengatter, T. (2005). A review of intervention studies aimed at household energy conservation. *Journal of Environmental Psychology*, *25*(3), 273–291.

ACM. (2019). *Energiemonitor 2019: Consumentenmarkt elektriciteit en gas.* Den Haag: Autoriteit Consument & Markt.

Aksoezen, M., Daniel, M., Hassler, U., & Kohler, N. (2015). Building age as an indicator for energy consumption. *Energy and Buildings*, *87*, 74–86.

Angrist, J. D., Imbens, G. W., & Rubin, D. B. (1996). Identification of Causal Effects Using Instrumental Variables. *Journal of the American Statistical Association*, *91*(434), 444–455.

Angrist, J. D., & Pischke, J. S. (2014). *Mastering 'metrics: The path from cause to effect*. *Mastering 'Metrics: The Path from Cause to Effect*. Princeton University Press.

Asafu-Adjaye, J. (2000). The relationship between energy consumption, energy prices and economic growth: Time series evidence from Asian developing countries. *Energy Economics*, *22*(6), 615–625.

Bandura, A. (1969). Theory Of ldentificatory Processes. *Handbook of Socialization Theory and Research.*, 213–262.

Behrangrad, M. (2015). A review of demand side management business models in the electricity market. *Renewable and Sustainable Energy Reviews*, *47*, 270-283.

Bertrand, M., Duflo, E., & Mullainathan, S. (2004). How much should we trust differences-in-differences estimates? *Quarterly Journal of Economics*, *119*(1), 249-275.

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

Burchell, K., Rettie, R., & Roberts, T. C. (2016). Householder engagement with energy consumption feedback: The role of community action and communications. *Energy Policy*, *88*, 168–177.

CBS. (2020). *CBS Open Data Portal*. Retrieved from CBS Statline: https://opendata.cbs.nl/statline/portal.html?_la=nl&_catalog=CBS

Card, D., & Krueger, A. B. (1994). Minimum wages and employment: a case study of the fast-food industry in New Jersey and Pennsylvania. *American Economic Review*, *84*(4), 772–793.

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

Cook, S. W., & Berrenberg, J. L. (1981). Approaches to Encouraging Conservation Behavior: A Review and Conceptual Framework. *Journal of Social Issues*, *37*(2), 73–107.

Darby, S. (2006). *A Review for DEFRA of the Literature on Metering, Billing and Direct Displays*. *Environmental Change Institute, University of Oxford*.

Darby, Sarah. (2010). Smart metering: What potential for householder engagement? *Building Research and Information*, *38*(5), 442–457.

De Lange, R. (2018). *Marktbarometer Aanbieding Slimme Meters: Voortgangsrapportage 2017.* Den Haag: Rijksdienst voor Ondernemend Nederland.

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

Depuru, S. S. S. R., Wang, L., & Devabhaktuni, V. (2011). Smart meters for power grid: Challenges, issues, advantages and status. *Renewable and Sustainable Energy Reviews*.

Enexis. (2020, Sept). *Planningsoverzicht Slimme Meter*. Retrieved from enexis.nl: www.enexis.nl/consument/aansluiting-en-meter/meter/slimme-meter/wanneer-een-slimme-meter/wanneer-komt-de-slimme-meter/planningsoverzicht

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

Franke, R. H., & Kaul, J. D. (1978). The Hawthorne Experiments: First Statistical Interpretation. *American Sociological Review*, *43*(5), 623.

Georgilakis, P. S. (2011). Environmental cost of distribution transformer losses. *Applied Energy*, *88*(9), 3146–3155.

Gillingham, K., Deng, H., Wiser, R., Darghouth, N. R., Nemet, G., Barbose, G., … Dong, C. (2016). Deconstructing solar photovoltaic pricing: The role of market structure, technology, and policy. *Energy Journal*, *37*(3), 231–250.

Gillingham, K., Rapson, D., & Wagner, G. (2016). The rebound effect and energy efficiency policy. *Review of Environmental Economics and Policy*, *10*(1), 68–88.

Grove, R. (2002). Climatic fears: Colonialism and the history of environmentalism. *Harvard International Review, 23*(4).

Haq, G., & Paul, A. (2013). *Environmentalism since 1945*. *Environmentalism since 1945*.

Hart, M., & De Dear, R. (2004). Weather sensitivity in household appliance energy end-use. *Energy and Buildings*, *36*(2), 161–174.

Heckman, J. J., Lalonde, R. J., & Smith, J. A. (1999). The economics and econometrics of active labor market programs. In *Handbook of Labor Economics* (Vol. 3 PART, pp. 1865–2097).

Hieminga, G. (2013). *Saving Energy in the Netherlands.* Amsterdam: ING Economics Department.

Hutton, R. B., Mauser, G. A., Filiatrault, P., & Ahtola, O. T. (1986). Effects of Cost-Related Feedback on Consumer Knowledge and Consumption Behavior: A Field Experimental Approach. *Journal of Consumer Research*, *13*(3), 327.

Kamp, H. (2014). *Voorzienings- en leveringszekerheid energie; Lijst van vragen en antwoorden over het ontwerp-besluit grootschalige uitrol slimme meters.* Den Haag: Parlementaire Monitor.

Kaza, N. (2010). Understanding the spectrum of residential energy consumption: A quantile regression approach. *Energy Policy*, *38*(11), 6574–6585.

Khandker, S.R.; Koolwal, G.B.; Samad, H. A. (2009). *Handbook on impact evaluation: quantitative methods and practices*. *The World Bank*. Washington D.C.

Lechner, M. (2010). The estimation of causal effects by difference-in-difference methods. *Foundations and Trends in Econometrics*.

Liander. (2020). *Open Data*. Retrieved from Liander: https://www.liander.nl/partners/datadiensten/open-data/data

Lin, J., Magnago, F. H., Foruzan, E., & Albarracín-Sánchez, R. (2017). Market Design Issues of Distributed Generation. In *Distributed Generation Systems: Design, Operation and Grid Integration* (pp. 369–413).

Lowe, P. (1985). The roots of modern environmentalism. *Journal of Rural Studies*, *1*(3), 294–295.

Mattern, F., Staake, T., & Weiss, M. (2010). ICT for green - How computers can help us to conserve energy. In *Proceedings of the e-Energy 2010 - 1st Int'l Conf. on Energy-Efficient Computing and Networking* (pp. 1–10).

Mayo, E. (2004). *The Human Problems of an Industrial Civilization*. *The Human Problems of an Industrial Civilization*.

Miller, G., Galanter, E., & Pribram, K. (1960). *Plans and the structure of behavior*. *Inc., New York*.

Moulton, B. R. (1990). An Illustration of a Pitfall in Estimating the Effects of Aggregate Variables on Micro Units. *The Review of Economics and Statistics*, *72*(2), 334.

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

Netbeheer Nederland. (2020, Aug). *Consumenteninformatie Slimme meter*. Retrieved from Netbeheer Nederland: www.netbeheernederland.nl/consumenteninformatie/slimmemeter

Nielsen, L. (1993). How to get the birds in the bush into your hand. Results from a Danish research project on electricity savings. *Energy Policy*, *21*(11), 1133–1144.

Parag, Y., & Sovacool, B. K. (2016). Electricity market design for the prosumer era. *Nature Energy*, *1*(4).

Pothitou, M., Hanna, R. F., & Chalvatzis, K. J. (2016). Environmental knowledge, pro-environmental behaviour and energy savings in households: An empirical study. *Applied Energy*, *184*, 1217–1229.

Pöttering, H., & Erlandsson, E. (2009). Directive 2009/72/EC of the European Parliament Concerning Common Rules for the Internal Market in Electricity. *Official Journal of the European Union*.

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

Rijkswaterstaat. (2020). *Klimaatmonitor Database*. Retrieved from Klimaatmonitor: https://klimaatmonitor.databank.nl/Jive

RVO. (2018). *Marktbarometer Aanbieding Slimme Meters - Voortgangsrapportage 2018.* Den Haag: Rijksdienst voor Ondernemend Nederland.

Schleich, J., Faure, C., & Klobasa, M. (2017). Persistence of the effects of providing feedback alongside smart metering devices on household electricity demand. *Energy Policy*, *107*, 225–233.

Sgouras, K. I., Bouhouras, A. S., Gkaidatzis, P. A., Doukas, D. I., & Labridis, D. P. (2017). Impact of reverse power flow on the optimal distributed generation placement problem. *IET Generation, Transmission and Distribution*, *11*(18), 4626–4632.

Stedin. (2020). *Verbruiksgegevens*. Retrieved from Stedin.net: www.stedin.net/zakelijk/open-data/verbruiksgegevens

Stock, H.J., & Watson, M. (2010). *Introduction to Econometrics (3rd Edition) (Addison-Wesley Series in Economics)*. *Addison-Wesley* (Vol. 1).

United Nations Development Programme. (2000). Energy, The Environment, and Health. In: *World Energy Assessment - Energy and the Challenge of Sustainability*.

Van Gerwen, R., Koenis, F., Schrijner, M., & Widdershoven, G. (2010). *Intelligente meters in Nederland; herziene financiële analyse en adviezen voor beleid.* Arnhem: KEMA Nederland B.V.

Van Wezel, B., & Van der Bie, R. (2015). *Elektriciteit in Nederland.* Den Haag: Centraal Bureau voor de Statistiek.

Van Wijnen, J. (2020, May 4). Netbeheerders vrezen voor fiasco met slimme meter. *Het Financieele Dagblad*, p. 1.

Vringer, K., & Dassen, T. (2016). *De Slimme Meter, Uitgelezen Energie(k)?* Den Haag: Planbureau voor de Leefomgeving.

Wilhite, H., & Ling, R. (1995). Measured energy savings from a more informative energy bill. *Energy and Buildings*, *22*(2), 145–155.

Winett, R. A., & Kagel, J. H. (1984). Effects of Information Presentation Format on Resource Use in Field Studies. *Journal of Consumer Research*, *11*(2), 655.

Wood, G., & Newborough, M. (2003). Dynamic energy-consumption indicators for domestic appliances: Environment, behaviour and design. *Energy and Buildings*, *35*(8), 821–841.

Wüstenhagen, R., & Bilharz, M. (2006). Green energy market development in Germany: effective public policy and emerging customer demand. *Energy Policy*, *34*(13), 1681–1696.

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

Yohanis, Yigzaw Goshu. (2012). Domestic energy use and householders' energy behaviour. *Energy Policy*, *41*, 654–665.

Zheng, J., Gao, D. W., & Lin, L. (2013). Smart meters in smart grid: An overview. In *IEEE Green Technologies Conference* (pp. 57–64).

Zipperer, A., Aloise-Young, P. A., Suryanarayanan, S., Roche, R., Earle, L., Christensen, D., Zimmerle, D. (2013). Electric energy management in the smart home: Perspectives on enabling technologies and consumer behavior. *Proceedings of the IEEE*.

Zuur, A. F., Ieno, E. N., & Elphick, C. S. (2010). A protocol for data exploration to avoid common statistical problems. *Methods in Ecology and Evolution*, *1*(1), 3–14.

Appendix A – Example of an Indicative Overview of User Cost

Explanation: the left-hand side of the overview contains the bi-monthly feedback on the electricity use

1. **Electricity use** over the past period, including the estimated user cost;
2. **Normative comparison** with the average consumption of a comparable user living in the same type of residency;
3. **Historic comparison** with the electricity use in previously recorded periods;
4. **Comparison for dummies**: you have done well.

Appendix B - Yearly Distribution in the Observations on Average Electricity Use

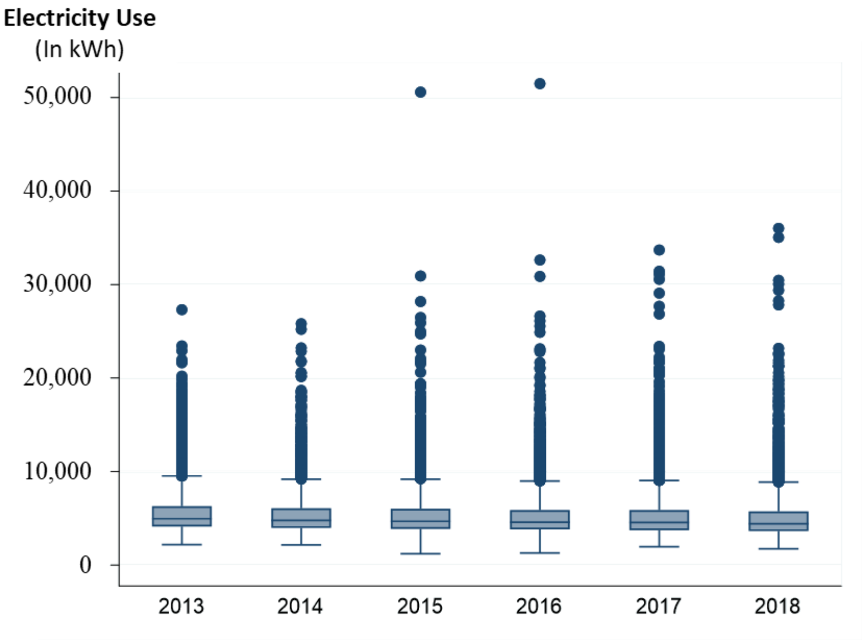

Appendix C – Yearly Distributions in the Percentage of Smart Meters per Postal Code Area

The values on the X-axis denote the proportion of smart meter connections within a postal code, the values on the Y-axis denote the frequency of observations for that particular proportion.

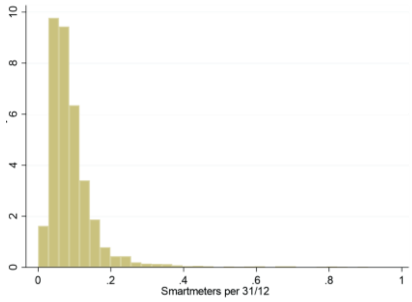

2013

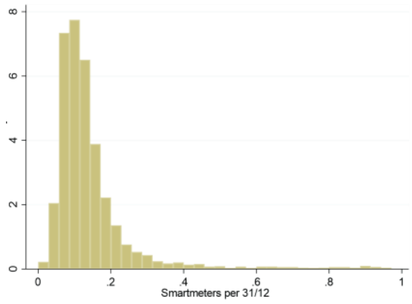

2014

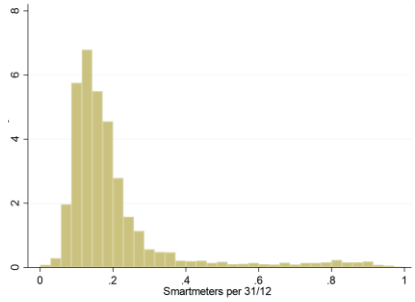

2015

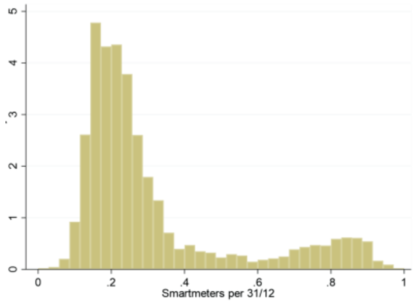

2016

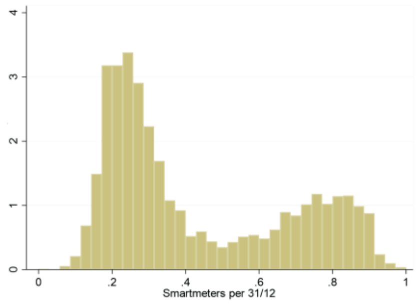

2017

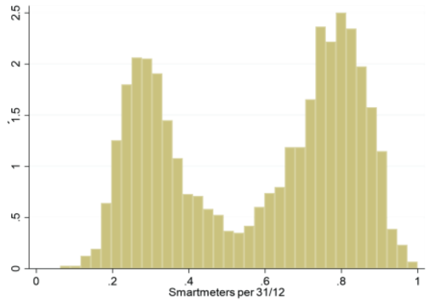

2018

Appendix D – Distribution of Observations: Before vs After adapting the Sample.

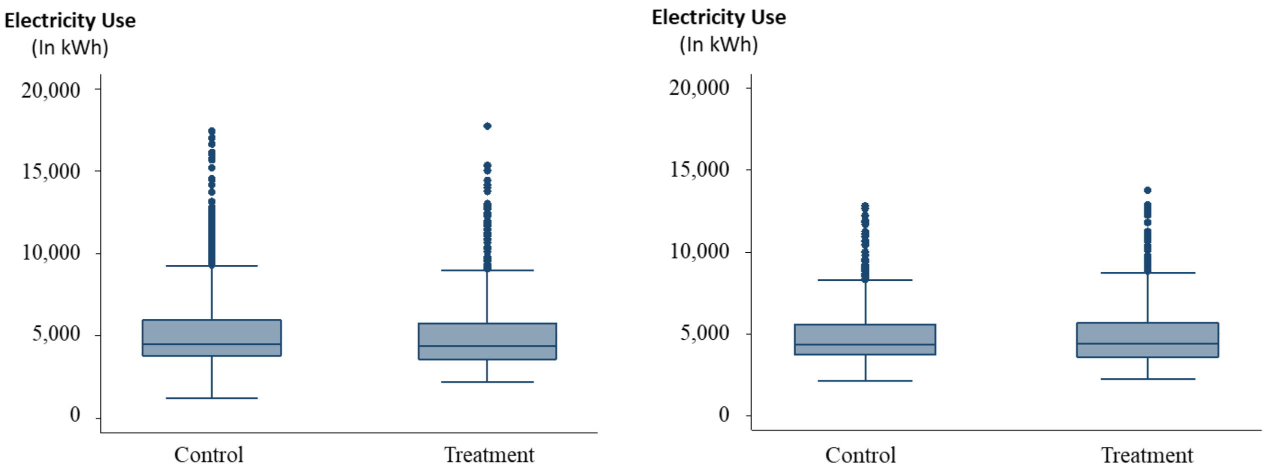

Appendix E – Robustness Check

TABLE 5
Estimation Results of the Fixed Effects Model

Connections	
- Smart Meter (%)	- 38.237
- Avg. Delivery (%)	1,188.353**
- Low Tariff Connections (%)	- 135.111
- District Heating (%)	15.196
- Gas Consumption (m^3)	0.316**
Population	
- Male (%)	- 2,523.895
- Dutch (%)	- 284.538
- Active (%)	529.028
Households	
- Household Size	305.631
- Rental (%)	- 251.000
- MDUs (%)	593.729*
- Avg. Income (€ 1,000)	10.643
Avg. Living Space	
- SDU's (m^2)	- 1.272*
- MDU's (m^2)	1.616*
Other	
- Charging Stations	- 1.586*
- EV's Registered	- 0.299
- Density	0.009
- Businesses	0.096
Interaction	40.605
Net Manager	Yes
Age Categories	Yes
Construction categories	Yes
Energy Labels	Yes
Area Fixed Effects	Yes
Time Fixed Effects	Yes
Observations	2,002
R^2	0.362

Notes: * denotes significance at the 10 % level; ** significant at the 5 % level; *** significant at the 1 % level. Robust standard errors clustered by postal code area (339 clusters) are included in the analysis, but not displayed in this Table.