Essays in Applied Microeconomics

Inauguraldissertation

zur Erlangung des Grades eines Doktors
der Wirtschaftswissenschaften

durch

die Rechts- und Staatswissenschaftliche Fakultät der
Rheinischen Friedrich-Wilhelms-Universität Bonn

vorgelegt von

Valentin Lorenz Stumpe

aus Königslutter am Elm

Bonn
2022

Dekan:	Prof. Dr. Jürgen von Hagen
Erstreferent:	Prof. Dr. Lorenz Götte
Zweitreferent:	Prof. Dr. Thomas Dohmen
Tag der mündlichen Prüfung:	29. April 2022

Acknowledgements

This thesis would not have been possible without the support of many great people. First, I would like to thank my two advisors, Lorenz Goette and Thomas Dohmen. Lorenz has been my supervisor during my Bachelor's thesis, my Master's thesis and during my PhD. Not growing tired of me (when there were plenty of opportunities) is something I give him high credit for. He supported me when I ran into problems, he provided valuable input when I was stuck, and he was relaxed when I was worried. For balancing the confidence he put in me with calming me down with his relaxed nature when things didn't go as I had planned, I will always be grateful. On this note, I would like to extend my thanks to Verena Tiefenbeck who, all those years ago during my Bachelor's, recommended asking him to supervise my Bachelor's thesis.

Thomas Dohmen was the best second (and definitely not second best!) supervisor I could have imagined. Finding the time to have regular meetings with his students, reading paper drafts and giving so much helpful feedback must have been difficult with such a full time schedule, but he even found time to have a quick chat before and after every talk to check on me and his other PhD students. For his continued guidance, especially during the last months of writing my thesis, for his enthusiasm in my projects and for invaluable feedback I am very grateful.

I would also like to thank Amelie Schiprowski, who, when I had just started working on my dissertation and felt a bit lost, took me under her wings, gave me feedback on my ideas and showed me the way around the economic research environment. It was so helpful to have an additional "informal advisor", and I'm glad I visited her presentation at the IZA three years ago.

Holger Gerhardt not only provided TeXnical assistance for my dissertation. Without his help, this thesis would not have been possible. This is not just a figure of speech - my laptop died a week before I had planned to hand in my thesis, and he helped me to recover my files and gave me his laptop to finish it.

The BGSE, especially Silke Kinzig and Andrea Reykers, helped me tremendously by guiding me through the administrative processes so I could focus on my research. The Collaborative Research Center Transregio 224 provided a great environment to exchange with other researchers and I am grateful to them for funding one of my research projects.

In addition, I would like to thank Justus Preusser, who was an enormous help and a great friend ever since we went to Berkeley. I would also like to thank the friends I made at the BGSE, for keeping me sane during coffee breaks, for all the valuable feedback, and for many lost games of Geotastic.

I would like to thank my family for their emotional support and finally, my Bachelor *Leistungslerngruppe*, who encouraged me to apply at the BGSE in the first place: Thank you, Max, Stella, Thorin and Valerie! And thank you Maike, for enduring and supporting me!

Contents

List of Figures

List of Tables

Introduction

Most economic models use traceable, easy-to-interpret parameters to illustrate basic economic concepts. The use of models helps us to get a grasp of the underlying, complex processes, without having to worry where the parameters stem from. From a policy point-of-view, it is common to map model predictions based on parameters into policy recommendations. Consequently, to produce clear policy recommendations, we need sound estimates of the underlying statistics. This holds true especially if the traditional assumption of rational agents - the *Homo Economicus* - fails. How much attention do we pay to the costs of certain goods? How do we react to price changes? How do we react to different economic policies? To answer these questions, we need real-world analyses of human behavior in economic contexts.

By using real-world data and exploiting quasi-random variation, I estimate crucial parameters of economic models, such as the salience parameter and price sensitivities in electricity consumption and job seeker responses to different governance policies. This dissertation consists of three independent chapters that seek to fill several gaps in the economic literature.

The first two chapters are connected by an overarching question: Can we use the electricity price to incentivize electricity consumption when it is plentiful and electricity conservation when it is scarce? At the current stage, electricity cannot easily be stored on a large scale and thus has to be produced when it is needed. And while fossil fuels can be stored to be used when needed, most renewable energy sources, such as solar and wind energy, are not available at all times, with their availability fluctuating in patterns with limited predictability. So far, policymakers have been focusing on the supply side of this challenge: During consumption spikes, additional fuel plants generate the electricity that is needed. But for microeconomists, the more interesting question to ask is: Can we approach the challenge from the demand side?

In Chapter 1 (joint with Lorenz Goette), I show that the effectiveness of energy prices to control electricity demand crucially depends on how much attention households pay to their electricity costs. Using high-frequency household electricity consumption data from a field experiment in Zurich, Switzerland, we first show that providing households with smart meters and In-Home-Displays to monitor their electricity consumption reduces domestic energy consumption. By exploiting a nonlinearity in the Swiss energy pricing mechanism, we show that feedback provision increases households' energy price sensitivity by more than 40 percent. We conceptualize these findings using a framework for GMM estimation. We find that due to salience bias, households behave as if they perceive less than 70 percent of their actual electricity costs. Heterogeneity analyses show that the treatment effect of feedback provision on the price sensitivity is increasing in pre-treatment baseline energy consumption and In-Home-Display usage. We also observe low-education and low-income households to be stronger biased than their highly educated, high-income counterparts. Finally, we show that the welfare implications of IHD-introduction differ significantly from pure monetary gains.

In Chapter 2 (joint with Lorenz Goette), I investigate the price elasticity of demand for households, which self-selected into a real-time electricity pricing tariff. Providing an explicit estimate of households' price sensitivities in a real-time pricing tariff allows for an even more precise prediction of electricity grid loads. We use hourly household consumption data from 899 households provided by a utility that passes hourly changing wholesale electricity prices on to the consumers. Using hourly wind energy production in Germany as an instrument for the hourly electricity price, we find that households indeed react strongly to real-time pricing. Using different levels of fixed-effects, we find that households react more strongly to price variation over the day than across days, with implied intra- and inter-day price elasticities of -1.67 and $-.45$, respectively. This implies that, while households are highly price sensitive, they show even stronger demand responses when consumption adjustment requires electricity load shifting within the day than across several days. Our results indicate that significant load shifts over the day can be achieved by using time-varying energy prices, which opens the possibility to offset energy consumption and production fluctuations using monetary incentives. Price reactions are strongest when prices are low, indicating that households evaluate price changes on their relative size to the price.

The third chapter (joint with Amelie Schiprowski and Patrick Arni) discusses a different policy-relevant question: How restrictively should the Public Employment Service (PES) manage job seekers? In many countries, the PES micro-manages job seekers' search behavior. This policy aims to ensure a quick reintegration into employment. However, it ignores many additional dimensions of job search, such as job match quality and externalities of job search. We study the effects of increasing the job search autonomy of job seekers on unemployment and labor markets. We exploit a policy change in the Swiss canton Bern, which strongly reduced search requirements and abolished mandatory vacancy referrals. Using detailed administrative data, we find that the policy change led to a reduction and narrowing of average job search. We set up a difference-in-differences design to estimate effects on labor market outcomes. Our results show that the policy change increased the average duration of unemployment spells in Bern by about 8%, while increasing average re-employment earnings by about 3%. Moreover, we find that effects are most pronounced among job seekers whose effort is predicted to decrease more strongly due to the reform. We also observe that the size of the effect depends on labor market conditions, such as labor market tightness or competition from job seekers from the remaining cantons. Finally, we find evidence that job seekers at the other side of the border benefited from the decrease in effort provided by job seekers in Bern by finding jobs at a faster rate.

Chapter 1

Quantifying the Salience Bias of Electricity Using Smart Meter Data

Joint with Lorenz Goette

1.1 Introduction

Over the last decades, growing concerns regarding climate change, energy security and sustainability have altered the public view of energy policy. With increasing political potential, energy efficiency concepts are rapidly rising in policymakers' priority. A central aspect of most energy sustainability concepts is a shift from from fossil-fuel energy towards wind and solar energy. Production from these sources, however, is not as easily controllable over time as their fossil-counterpart. Hence, in order to push forward sustainable energy production, a major future task for both researchers and policymakers alike will be to find ways to incentivize electricity demand to follow its supply. One of the main instruments to control energy demand is the electricity price, which in turn can be manipulated through levies, taxes and subsidies. These seem like powerful tools: In Germany as well as other European countries, a large portion of the energy price consists of taxes and levies.[1] This provides a starting point for interventions from policymakers to control aggregate energy consumption by manipulating the price. It is tempting to look at energy prices and price elasticities the same way as we look at other consumption goods: Increasing the price should decrease demand and hence, consumption. However, in reality, electricity consumers exhibit a remarkably low price elasticity (Reiss and White (2005), Allcott (2011), Ito (2014)). The reason for this low elasticity is subject to an ongoing debate.

1. Source: Bundesnetzagentur, (https://www.bundesnetzagentur.de/DE/Beschlusskammern/BK08/BK8_06_Netzentgelte/BK8_NetzE_basepage.html). Last visited Jan 9, 2022.

In this paper, we explore what we believe to be one of the main underlying mechanisms determining domestic electricity consumption by quantifying how much attention households pay to electricity costs compared to more conspicuous, immediate costs. We present evidence from a randomized field experiment in Zurich, Switzerland and enhance the original design with a strategy to identify electricity price elasticities. In the original setting, 1'176 households were provided with smart meters, which recorded electricity consumption in 15-minute intervals, and In-Home-Displays (IHDs), which allowed households to receive continuous feedback about their current and past electricity consumption and to set personal consumption goals (Degen, Efferson, Frei, Goette, and Lalive (2013)). We extend the original experiment's empirical strategy by exploiting the nonlinear electricity pricing mechanism on weekends in Zurich to compare electricity price elasticities with and without real-time feedback, using households with an installed smart meter, but no feedback through IHD as a control group. We then impose a theoretical framework to quantify how much attention consumers pay to electricity costs before IHDs make them fully salient.

Following IHD-introduction we find a reduction of electricity consumption during the day by 0.316 kWh, or 5 to 6 percent of total daily consumption and a somewhat ambiguous effect on nighttime electricity consumption. Using a framework with a quadratic utility function and variation in the daytime electricity price for different days of the week, we find negative price elasticities for both daytime and nighttime electricity consumption with respect to the daytime electricity price. Building up on our results, we estimate that households only perceive 68 percent of the actual costs incurred from electricity consumption. We hypothesize that the introduction of IHDs makes electricity costs more salient and hence, incentivises households to shift energy consumption from high-costs days of the week to low-costs days of the week, while load shifting between different times of the day is subject to spillover-effects.

This paper contributes to the literature on the application of social and behavioral sciences to research on energy consumption (Abrahamse, Steg, Vlek, and Rothengatter (2007), Ehrhardt-Martinez (2008), Ehrhardt-Martinez and Laitner (2010), Ueno, Sano, Saeki, and Tsuji (2006), Wilson and Dowlatabadi (2007), Steg and Vlek (2009)) and shrouded attributes of resource consumption (Gabaix and Laibson (2006), Sunstein (2015)). More fundamentally, we expand the literature on price elasticities of energy consumption (Lijesen (2007), Al-Faris (2002), Aigner, Newman, and Tishler (1994), Filippini (1995), Patrick and Wolak (2001), Allcott (2011), Ito (2014), Jessoe and Rapson (2014), Reiss and White (2005)) by exploring the mechanisms behind the low price elasticities found in the literature.

Previous empirical work has shown in various contexts that electricity consumers exhibit a remarkably low price elasticity (Reiss and White (2005), Allcott (2011), Ito (2014)). There are two apparent reasons why consumer demand may be inelastic: first, the underlying preferences may be such that demand responses to price changes are indeed inelastic. In that case, taxes and subsidies on household electricity consumption would only lead to small changes in electricity consumption. On the other hand, there is also increasing evidence that that individuals' demand may be inelastic due to inattention to their electricity use and costs. This inattention bias towards the less salient at the time of decision making has been found in a variety of fields and is a second potential driver behind low price elasticities (Chetty, Looney, and Kroft (2009), Jessoe and Rapson (2014)).

In many settings, some features of a decision are vivid and salient, while others are less present and harder to quantify. This imbalance in the salience of some features of a consumption decision disproportionately engages our attention and hence, distorts our economic decision making, resulting in a salience bias. Electricity consumption is especially prone to the salience bias: The benefits from using energy (watching television, boiling water etc.) are felt immediately. The costs, however, are incurred with a delay and are thus not as salient. When turning on an electric device, consumers do not open their purse and pay so-and-so many cents per kilowatt-hour used. Instead, they receive a bill at the end of the month or at the end of the year. This makes the costs of electricity consumption less salient than they are in most economic contexts in which consumers pay upfront. A growing branch of the behavioral economics literature on resource consumption has thus focused on feedback provision as a means to close the salience gap between benefit and costs (Attari, DeKay, Davidson, and Bruine de Bruin (2010), Attari (2014), Allcott and Mullainathan (2010)). Indeed, feedback has been shown to reduce energy- and water consumption by making their respective costs more salient.

Why does feedback play such an important role for conservation efforts? In many ways, electricity differs from other consumption goods. It is no visible product, neither when it is produced, nor when it is used. Consumers do not see a “diminishing stock” of electricity when turning on the television, nor do they have a precise idea how the energy they are using was produced. Not being able to directly perceive the financial implications makes it difficult for consumers to develop awareness or involve in such an abstract concept. This yields special economic mechanisms for electricity markets: For example, as electricity is no directly visible good, branding mechanisms work differently (Hortaçsu, Madanizadeh, and Puller (2017), Rutter, Chalvatzis, Roper, and Lettice (2018)). This makes it harder for green energy to be a “lifestyle-element” that can be shown around, such as organic food or eco-clothing (Birzle-Harder and Götz (2001)). Additionally, in most developed countries, electricity shortages usually do not occur. Hence, consumers see electricity as a seemingly unlimited, barely tangible resource, which they do not consume directly, but rather as an input for electronic devices. The inconspicuous nature of household electricity consumption thus rarely begs questions regarding its origin, usage or costs, especially since electricity costs only make up a small share of household expenditures.

While information provision strategies prove to be effective tools for consumption reduction in many environmentally relevant contexts, feedback appears to be among the most effective. Fischer (2008) points out that successful feedback features high frequency, long runtime and appliance-specific breakdown, is presented in a clear way and uses a certain degree of interaction with the consumer. The most appealing way to satisfy these requirements seems to be smart metering and the use of In-Home-Displays (IHDs) to enable households to continuously monitor their electricity consumption. This can potentially avoid the cyclical pattern of "action-and-backsliding" found by Allcott and Rogers (2014), where strong reactions to feedback provision are quickly followed by a return to old behavioral patterns following infrequent electricity feedback. In some cases, the IHDs are able to break down the feedback on appliance level, enabling consumers to see exactly which appliance uses how much energy, or provide normative feedback in the form of social comparisons (Andor, Gerster, and Goette (2020)).

The remainder of this paper is structured as follows: In Section 1.2, we develop a simple model of domestic electricity consumption. Section 1.3 provides an overview of the field experiment conducted in Zurich, Switzerland. In Section 1.4, we describe our methodology to identify the framework developed in Section 1.2 and present the results. Section 1.5 presents the welfare implications of our findings. Section 1.6 concludes.

1.2 The Theoretical Framework

In this section we introduce a model of household energy consumption to illustrate and discuss the underlying parameters and mechanisms. Our model simplifies and loosely builds up on several concepts introduced by DellaVigna (2009) and Chetty, Looney, and Kroft (2009). Using the concepts introduced here, we will track the components that determine a household's energy consumption decision to motivate our empirical strategy. Later, we will use this framework to quantify the salience bias of the average household. In order to do so, we will make stricter assumptions on the form of the utility function. For now, we will keep the framework as simple and general as possible: A household can consume different quantities of two distinct goods: Daytime electricity consumption and nighttime electricity consumption. In most countries, electricity consumption during the day costs more than at night.[2] Hence, we can treat daytime and nighttime electricity consumption as different goods with different prices to investigate the underlying price and substitution effects.

1.2.1 Basic Setup

Let a household's utility be determined by its consumption of electricity at different times of the day. To keep the framework simple and to match it to the existing

2. We will discuss the reasons for this pricing mechanism in Section 1.4.

pricing mechanism, we will only distinguish electricity consumption during the day (denoted as x_1) and during the night (denoted as x_2).
A household's utility is defined as

$$U(x_1, x_2, p_1, p_2) \;=\; u(x_1, x_2) - \theta(p_1 x_1 + p_2 x_2)$$

Where u denotes the utility from consuming electricity (e.g. watching TV or showering). The parameter $\theta < 1$ is the household's level of attention to energy costs, as in Chetty, Looney, and Kroft (2009) or DellaVigna (2009). The intuition in our research context is that while the utility from electricity consumption is felt immediately and correctly, the associated resource use is difficult to perceive. As in Chetty, Looney, and Kroft (2009) and Chetty (2009), we assume that individuals only give weight θ to the energy costs due to limited attention. As $\theta \to 1$, the quantity, and hence, the cost of consumption is correctly perceived.

This specification for attention in a demand model is reduced-form in the sense that it does not specify a deeper micro foundation. A plausible interpretation, along the lines of Enke and Graeber (2019), is that individuals need to pay attention to perceive their true electricity use: they observe a signal $z = x + u$, where $x \sim N(x^D, \sigma_x^2)$ is the distribution of their perceived electricity use, and $u \sim N(0, \sigma_{u,t}^2)$ is a perception error due to limited attention that is centered around zero. Given a signal z, the individual rationally infers that her electricity use x is

$$E(x|z) \;=\; \underbrace{\theta x + \theta u}_{\equiv \theta z} + (1-\theta)x^D.$$

Thus, the attention parameter can be thought of as the signal-to-noise ratio $\theta = \frac{\sigma_x^2}{\sigma_x^2 + \sigma_{u,t}^2}$, arising from this signal-extraction problem under limited attention.[3] We make the following assumptions:

Assumption 1. (Assumptions on u, θ and p):
- (i) $0 < \theta < 1$
- (ii) $\frac{\partial u}{\partial x_i} = u_i > 0,\ i = 1, 2$
- (iii) $\frac{\partial^2 u}{\partial x_i^2} = u_{ii} < 0,\ i = 1, 2$
- (iv) $\frac{\partial^2 u}{\partial x_1 \partial x_2} = u_{ij} < 0$
- (v) $u_{11} - \frac{u_{12}^2}{u_{22}} < 0$
- (vi) $p_1 > p_2$

3. Tiefenbeck, Goette, Degen, Tasic, Fleisch, et al. (2016) show that, in the context of water consumption, individuals with below-average water use tend to over-estimate their water use, while individuals with above-average water use tend to underestimated their water use.

The first two conditions are self-explanatory: Consumers are salience-biased, and utility should be increasing in energy consumption. Note that (iii) and (iv) ensure constant and negative second derivatives, and (vi) assumes nighttime electricity to be cheaper than daytime electricity. Taking derivatives, the first-order conditions deliver:

$$u_1(x_1^*, x_2^*) = \theta p_1 \tag{1.1}$$

$$u_2(x_1^*, x_2^*) = \theta p_2 \tag{1.2}$$

This induces the following comparative statics:

$$\begin{aligned}
\frac{\partial x_1}{\partial \theta} &= \frac{p_1 - \frac{u_{12}}{u_{22}} p_2}{u_{11} - \frac{u_{12}^2}{u_{22}}} \\
\frac{\partial x_2}{\partial \theta} &= \frac{p_2 - \frac{u_{12}}{u_{11}} p_1}{u_{22} - \frac{u_{12}^2}{u_{11}}} \\
\frac{\partial x_1}{\partial p_1} &= \frac{\theta}{u_{11} - \frac{u_{12}^2}{u_{22}}} \\
\frac{\partial x_2}{\partial p_2} &= \frac{\theta}{u_{22} - \frac{u_{12}^2}{u_{11}}} \\
\frac{\partial x_1}{\partial p_2} &= -\frac{u_{12}}{u_{11}} \frac{\partial x_2}{\partial p_2} \\
\frac{\partial x_2}{\partial p_1} &= -\frac{u_{12}}{u_{22}} \frac{\partial x_1}{\partial p_1}
\end{aligned}$$

In particular, the effect of correcting the salience bias (exogenously manipulating θ) in daytime and nighttime electricity consumption depends on the respective prices p_1, p_2 and substitution effects. We can also express the treatment effect as a function of the salience parameter θ and the price effects:

$$\frac{\partial x_1}{\partial \theta} = \frac{1}{\theta}\left[p_1 \frac{\partial x_1^*}{\partial p_1} + p_2 \frac{\partial x_1^*}{\partial p_2}\right] \tag{1.3}$$

$$\frac{\partial x_2}{\partial \theta} = \frac{1}{\theta}\left[p_2 \frac{\partial x_2^*}{\partial p_2} + p_1 \frac{\partial x_1^*}{\partial p_2}\right] \tag{1.4}$$

1.2.2 Introducing Smart Metering

Introducing smart meters and In-Home-Displays allows households to receive real-time feedback on their energy consumption and thus, provides variation in the salience parameter θ.

Proposition 1. (Consumption effect of IHD-introduction):
Manipulating θ can have positive or negative effects on electricity consumption:

$$\frac{\partial x_1}{\partial \theta}\begin{cases} < 0 & for\ p_1 > \frac{u_{12}}{u_{22}}p_2 \\ \geq 0 & otherwise \end{cases}$$

$$\frac{\partial x_2}{\partial \theta}\begin{cases} < 0 & for\ p_2 > \frac{u_{12}}{u_{11}}p_1 \\ \geq 0 & otherwise \end{cases}$$

The intuition behind this proposition is as follows: While the price effect of correcting the salience bias should decrease daytime and nighttime energy consumption, the substitution effect may work in a different direction: If the price of nighttime electricity is much lower than the price of daytime electricity ($p_2 \leq \frac{u_{12}}{u_{11}}p_1$), households will excessively reduce expensive daytime energy consumption and substitute it with cheaper nighttime energy consumption. Introducing IHDs would then increase x_2. This load shifting effect is regularly found in the literature.
In the following, we will assume that IHD-introduction fully resolves the salience bias, setting θ to 1. Note that we can easily relax this assumption and instead assume that providing households with feedback instead sets θ to some $\theta^{FB} < 1$. In that case we would identify the salience parameter as $\tilde{\theta} = \theta^{FB} \cdot \theta$, where θ denotes the salience parameter we would have identified if we had assumed feedback to fully resolve the salience bias.

1.2.3 Identifying the Salience Bias in a Concrete Framework

In order to identify the salience parameter θ outlined above, we have to make stronger assumptions on $u(x_1, x_2)$.
Let $u(x_1, x_2) = \alpha_1 x_1 + \alpha_2 x_2 + \gamma x_1 x_2 - \beta_1 x_1^2 - \beta_2 x_2^2$.
Additionally, we need a few regularity conditions on the model parameters:

Assumption 2. (Assumptions on the utility function parameters):
- (i) $|\gamma| < 2\sqrt{\beta_1 \beta_2}$
- (ii) $\alpha_1 > \alpha_2 > 0$
- (iii) $\beta_2 \geq \beta_1 > 0$

(i) ensures concavity of the utility function. (ii) and (iii) imply that daytime energy consumption generates higher utility than nighttime energy consumption.[4]
Again, taking derivatives yields the first-order conditions:

4. Of course, some implications of the quadratic utility function we chose (such as constant derivatives with respect to the salience parameter or electricity prices), may appear odd. To see this, consider two households that hugely differ in baseline electricity consumption: One household exhibits a baseline of 4 kWhs per day, the second household 10 kWhs per day. Increasing the electricity price should have a larger effect on the second household, as it exhibits higher incentives and a larger scope for consumption reduction. We do not consider this possibility. We will, however provide results based

$$x_1^* = \frac{1}{4\beta_1\beta_2 - \gamma^2}[2\beta_2(\alpha_1 - \theta p_1) + \gamma(\alpha_2 - \theta p_2)] \tag{1.5}$$

$$x_2^* = \frac{1}{4\beta_1\beta_2 - \gamma^2}[2\beta_1(\alpha_2 - \theta p_2) + \gamma(\alpha_1 - \theta p_1)] \tag{1.6}$$

This delivers the following comparative statics:

$$\begin{aligned}
\frac{\partial x_1^*}{\partial \theta} &= -\frac{1}{4\beta_1\beta_2 - \gamma^2}(2\beta_2 p_1 + \gamma p_2) \ < \ 0 \\
\frac{\partial x_2^*}{\partial \theta} &= -\frac{1}{4\beta_1\beta_2 - \gamma^2}(1\beta_1 p_2 + \gamma p_1) \\
\frac{\partial x_1^*}{\partial p_1} &= \psi_{11}^* \ = -\frac{1}{4\beta_1\beta_2 - \gamma^2}(2\beta_2\theta) \ < \ 0 \\
\frac{\partial x_1^*}{\partial p_2} &= \frac{\partial x_2^*}{\partial p_1} \ = \psi_{21}^* \ = -\frac{1}{4\beta_1\beta_2 - \gamma^2}(\gamma\theta)
\end{aligned} \tag{1.7}$$

Proposition 2. (Constant Derivatives)
Assuming a quadratic utility function, the derivatives of x_1 and x_2 with respect to their prices and the salience parameter θ are constant in x_1 and x_2.

In order to identify θ, we will define the price- and cross-elasticities for daytime electricity consumption (i.e. x_1) under feedback. Note that we can do this analogously for x_2 (nighttime electricity consumption). However, in the data, our price variation is limited to the daytime electricity price, which is why we will restrict our theoretical analysis to identifying θ using price variation in daytime electricity. We now define the elasticities:

$$\begin{aligned}
\epsilon_1^{FB} &= \psi_{11}^{FB}\left(\frac{p_1}{x_1}\right) = \frac{2\beta_2}{H}\left(\frac{p_1}{x_1}\right) \\
\phi_1^{FB} &= \psi_{21}^{FB}\left(\frac{p_2}{x_1}\right) = \frac{\gamma}{H}\left(\frac{p_1}{x_1}\right)
\end{aligned}$$

with $H = \gamma^2 - 4\beta_1\beta_2$.
Using our assumption that feedback fully resolves the salience bias and sets θ to 1, and that the derivative $\frac{\partial x_1^*}{\partial \theta}$ is constant in x_1, we have

$$\frac{\partial x_1^*}{\partial \theta} = \frac{\Delta x_1^*}{1 - \theta}$$

on baseline consumption tercile splits in Section 1.4.7. Additionally, one can see the utility function we chose as a second order approximation of the *true* utility function.

Where Δx_1^* denotes the change in daytime electricity consumption following the introduction of IHDs, and $1-\theta$ denotes the increase in salience following IHD-introduction.
Then, equation (1.7) delivers

$$\theta = 1 - \frac{\Delta x_1^*}{p_1\psi_{11}^{FB} + p_2\psi_{12}^{FB}} = 1 - \frac{\Delta x_1^*}{p_1\psi_{11}^{FB} + p_2\psi_{21}^{FB}} \tag{1.8}$$

Looking at our theoretical framework, we can find an alternative method to identify the salience parameter θ. To see this, consider the expression for ψ_{11}^* from Section 1.2.3 and let ψ_{11}^{FB} denote the derivative of daytime electricity consumption with respect to the daytime electricity price. Then:

$$\psi_{11}^* = -\frac{1}{4\beta_1\beta_2-\gamma^2}(2\beta_2\theta) \tag{1.9}$$

$$\psi_{11}^{FB} = -\frac{1}{4\beta_1\beta_2-\gamma^2}(2\beta_2) \tag{1.10}$$

Naturally, this allows us to extract θ as the ratio of the two expressions:

$$\theta = \frac{\psi_{11}^*}{\psi_{11}^{FB}} = \frac{\Delta x_1^{*,p_1}}{\Delta x_1^{FB,p_1}} \tag{1.11}$$

Where $\Delta x_1^{*,p_1}$ denotes the change in pre-intervention DT-consumption following a change in p_1 and the second equality uses our assumption of constant derivatives. Since Δp_1 does not change over time, it cancels out in the expression above. Hence, the inattention parameter can be identified as the ratio of the price sensitivities before and after IHD-installation. Intuitively, the more inattentive a consumer is towards electricity costs (pre-intervention), the smaller the pre-intervention price sensitivity should be relative to the post-intervention price sensitivity under full attention.

1.2.4 Consumer Surplus Analysis

In our framework, households are salience-biased. With $\theta < 1$, a household does not optimize with respect to the consumer surplus, but underestimates the costs of electricity consumption.
A household's *unbiased utility*, i.e. the consumer surplus, is defined as

$$W(x_1,x_2,p_1,p_2) = u(x_1,x_2) - p_1x_1 - p_2x_2$$

With $u(x_1,x_2)$ defined as above. Plugging in the optimal choices of x_1 and x_2 under salience bias from (1.5) and (1.6) and deriving with respect to θ gives us the surplus effect of introducing smart metering and IHDs and thus, correcting the salience bias.[5]

$$\frac{\partial W(\theta)}{\partial \theta} = -(1-\theta)\left(p_1\frac{\partial x_1^*}{\partial \theta} + p_2\frac{\partial x_2^*}{\partial \theta}\right) \tag{1.12}$$

Integrating with respect to θ and taking differences gives us:

$$\begin{aligned} W(1) - W(\theta) &= -\left(p_1\frac{\partial x_1^*}{\partial \theta} + p_2\frac{\partial x_2^*}{\partial \theta}\right)\int_\theta^1 (1-s)ds \\ &= -\frac{1}{2}(1-\theta)(p_1\Delta x_1^* + p_2\Delta x_2^*) \end{aligned} \tag{1.13}$$

Plugging in our expression for θ from equation (1.8) delivers:

$$W(1) - W(\theta) = -\Delta x_1^* \frac{(p_1\Delta x_1^* + p_2\Delta x_2^*)}{2\left(p_1\psi_{11}^{FB} + p_2\psi_{21}^{FB}\right)} \tag{1.14}$$

Proposition 3. (Surplus effects of IHD-introduction)
The surplus effect of resolving the salience bias does not simply equal the savings from using less electricity, but is scaled by the substitution effects of daytime- and nighttime consumption and is dampened by the fact that consuming less energy also decreases utility.

1.2.5 The Role of Heterogeneity in Consumer Surplus Gains

We can also express the surplus effect of IHD-introduction in terms of the salience parameter θ and the electricity price effects. Using equations (1.3) and (1.4), we can write

$$\frac{\partial W(\theta)}{\partial \theta} = -(1-\theta)\left[p_1^2\frac{\partial x_1^{FB}}{\partial p_1} + 2p_1p_2\frac{\partial x_1^{FB}}{\partial p_2} + p_2^2\frac{\partial x_2^{FB}}{\partial p_2}\right]$$

And so

$$W(1) - W(\theta) = \frac{1}{2}p'Hp \cdot (1-\theta)^2 \tag{1.15}$$

with $p' = (p_1, p_2)$ and $H = \begin{pmatrix} \frac{\partial x_1^{FB}}{\partial p_1} & \frac{\partial x_1^{FB}}{\partial p_2} \\ \frac{\partial x_1^{FB}}{\partial p_2} & \frac{\partial x_2^{FB}}{\partial p_2} \end{pmatrix}$

5. For clarity reasons we will change notation such that $W(x_1^{FB}, x_2^{FB}, p_1, p_2) = W(1)$ and $W(x_1^*, x_2^*, p_1, p_2) = W(\theta)$

Taking expectations yields

$$E[\Delta W] = \frac{1}{2}p'Hp \cdot E[(1-\theta)^2]$$
$$= \frac{1}{2}p'Hp \cdot E\left[(E[1-\theta])^2 + V(1-\theta)\right]$$

Proposition 4. (The role of heterogeneity)
For a given mean bias, heterogeneity in the salience parameter θ increases the welfare gains from resolving the salience bias.

1.3 The EWZ Field Experiment

In this section, we present the randomized field experiment conducted by Degen et al. (2013) and augment their strategy with our empirical approach. 85.955 randomly selected costumers of the electric power company in Zurich (EWZ) were contacted and invited to participate in a study. If they were interested, they were asked to register online and give their consent for a data protection statement. Households were then randomly selected into one of five intervention groups and 26 starting cohorts[6], respectively. Of the total study population, 1.176 households (out of which, 1'105 households completed the study and 920 delivered usable data and completed all surveys throughout the study) were provided with a smart meter device and an In-Home-Display. Hence, their electricity consumption could be recorded and households received continuous feedback on their electricity consumption. The smart meter recorded electricity consumption in 15-minute intervals for several weeks before the IHD was installed (to document baseline consumption) and more than one year after installation of the IHD. After IHD-installation, households could, at any time, see their current and past electricity consumption. For a more detailed illustration of the IHD-interface, see Appendix 1.A. Before presenting our empirical approach, we will take a short look at the original implementation of the field experiment, the treatment interventions, the data structure and descriptive statistics.

1.3.1 Setting and Implementation

The city of Zurich is the capital of the Swiss canton with the same name, with a population of roughly 372.000 in 2010, when the experiment started.[7] The implementation of the field experiment followed three steps:

6. Each cohort started the study at a different point in time as delivery and installation of the smart meters and In-Home-Displays could not be managed for all households at the same time.

7. Source: Department of Statistics, Canton Zurich (https://statistik.zh.ch, last visited: Jan 9, 2022).

1. Recruitment of households: Households were mailed invitations to participate in the study.
2. Selection of households and first survey: After registration, out of the households which were willing to participate, 5.000 were selected and sent a first survey regarding the household's demographics.
3. Assignment to treatment groups: Participants were assigned to one of five experimental groups. The objective of the assignment algorithm was to minimize the variance of baseline energy consumption between the experimental groups. The experimental groups were:
   - G0 Control Group: No intervention was implemented here. Note that for our empirical analysis, we will use a different control group, consisting of households with a smart meter, but no IHD.[8]
   - G1 Smart Metering Group: This group of households received a smart meter and, after several weeks, an In-Home-Display to continuously receive feedback on their energy consumption. Households in this group were also randomly selected into squadrons, each starting the study at a different time, due to delivery constraints with the smart meters and IHDs. We will focus our analysis on this treatment group.
   - G2 Consulting Group: Households in this treatment group received professional advice on electricity conservation.
   - G3 Social Competition Group: Each household in this treatment group was assigned to a similar household in terms of baseline energy consumption and household characteristics. The households then received monthly, and later quarterly, feedback regarding their own and their partner's electricity consumption. Both parties knew about their partner's information set.
   - G4 Social Comparison Group: Each household was assigned a comparable household from the control group G0 and could compare their and their partner's electricity consumption. Feedback for this intervention group was one-directional, the partner household did not receive any information.

Figure 1.1 illustrates the timeline (in weeks) for G1-households. Rollout of the IHDs occurred between January 2011 and July 2012, and was set several weeks (in most cases eight weeks) after installation of the smart meters. This means that T does not refer to a specific point in time that is the same for all households. Instead, it denotes a squadron-specific date between January 2011 and July 2012. The assignment to these dates within the G1-group was random. Our consecutive analysis uses this fact and exploits that during a substantial part of the experiment, a large portion of the smart meters was already installed, but not yet linked to the IHDs, providing us

8. We will evaluate on this point in Section 1.4.

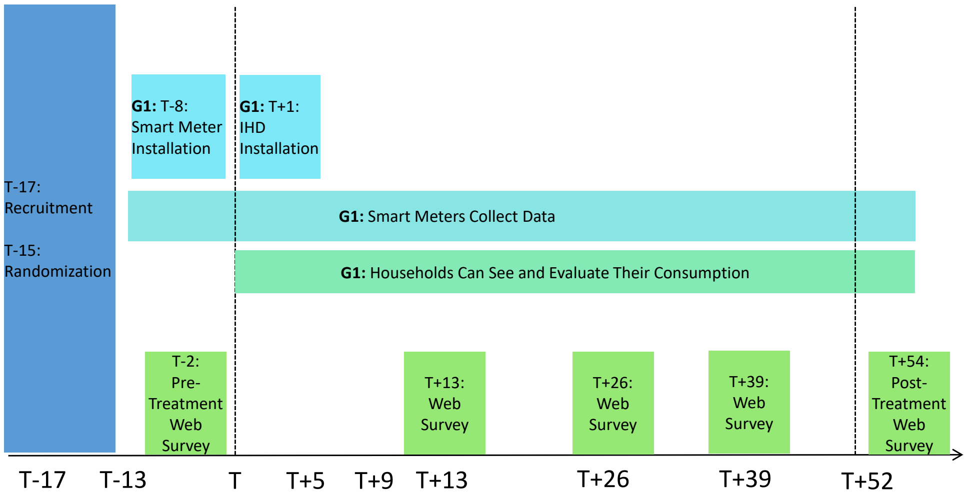

Figure 1.1. Timeline For G1-Households

This figure presents the timeline for households in the G1 Smart Metering Group. Households in this group received a smart meter several weeks before IHD-installation to collect information on baseline consumption. At this point in time, the smart meters started collecting consumption information on the households, but households did not receive real-time feedback until IHD-installation. Once the IHDs were installed (at time *T*), households could see and evaluate their electricity consumption continuously.

with an additional control group for the smart meter analysis: households with an installed smart meter but without IHD provide high-frequency data on electricity consumption for households *without* IHD, whereas the regular control group G0 does not provide high-frequency data. We will return to this point when presenting the graphical evidence and our empirical strategy in Section 1.4.

1.3.2 Data Overview

In the following, unless otherwise specified, we will exclusively focus on the Smart Metering Group G1. Our data contains not only information on electricity consumption in 15-minutes intervals, but also on IHD usage activity as well as survey data on household characteristics, such as number of household members, household income, education level, environmental attitudes etc.
The smart meters record 15-minute-interval meter readings for daytime and nighttime electricity consumption. Note that each household in Zurich has two electricity

readers: One records total daytime consumption, the other records total nighttime consumption since installation. This mechanism is used to distinguish daytime- and nighttime consumption when billing the household.

1.3.3 Summary Statistics

Table 1.1. Household Characteristics

	G0 Control	G1 Smart Meter	G2-G4 Other Treatments	Zurich
Demographics				
Female	0.362	0.370	0.380	0.504
Age[a]				
0-19	0.30	0.09	0.17	19.90
20-39	42.98	43.71	43.36	26.30
40-64	46.15	44.79	43.99	35.00
64+	10.57	11.40	12.49	18.80
Nationality				
Swiss	0.838	0.802	0.818	0.675
German	0.090	0.119	0.109	0.081
Italian	0.012	0.015	0.016	0.035
Serbian	0.002	0.003	0.001	0.008
Household Information				
Household Size	2.128	2.073	2.078	1.990
Tenancy[b]	0.888	0.894	0.903	0.900
Pre-treatment-per-day-consumption (kWh per day)	6.107	6.343	6.020	$\approx$ 7.534
N (Age)	1’012	1’105	2’994	372’000

Notes: This table presents household characteristics in our sample, compared to the average household in Zurich and Switzerland. (Source for the Zurich demographics: Swiss Federal Office of Statistics (https://www.bfs.admin.ch/bfs/de/home/statistiken/bevoelkerung/stand-entwicklung/bevoelkerung.html, last visited Dec 20, 2021), the Canton of Zurich (https://www.zh.ch/de/planen-bauen/raumplanung/immobilienmarkt/wohnungsmieten.html, last visited Jan 10, 2022), and Swiss Agency for Energy Efficiency)
[a]Numbers in our sample refer to the share of survey respondents falling into the respective age category
[b]0 = Homeowner/Condominium; 1 = Rented Apartment

Table 1.1 presents the household characteristics in our sample by experimental groups and compares them to official data from Zurich.[9] We first see that females are underrepresented in our sample. Swiss citizens are overrepresented in the study, as are Germans, while Italians and Serbs are underrepresented. This can mostly be attributed to the fact that the study was conducted in German. On average, a household in our sample has more members than the average household in Zurich, but the difference is negligible. Finally, note that the distribution of the age of survey respondents does not perfectly translate into Zurich demographics, as household heads in our sample, who are usually the surveys' respondents, are naturally usually older than 19 years and we do not have information on the age of the remaining household members. Figure 1.2 presents the distribution of baseline electricity consumption.

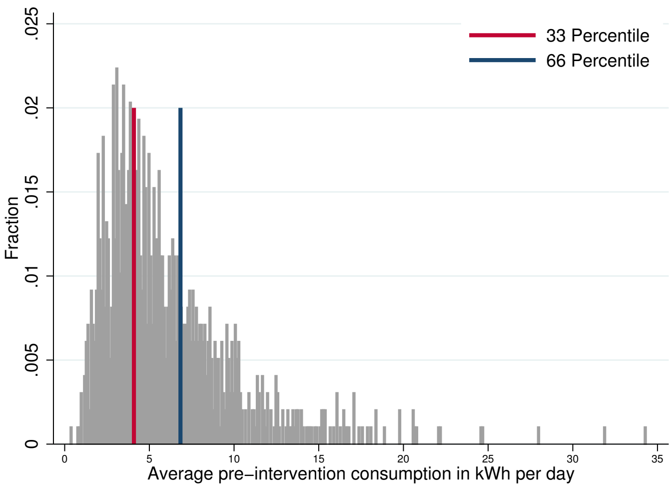

Figure 1.2. Distribution of Daily Baseline Consumption

This figure presents the distribution of average daily baseline consumption per household on Saturdays and Sundays during the weeks before IHD-installation. The red and blue line denote the 33 and 66-percentile of baseline electricity consumption, respectively.

1.4 The Empirical Analysis

In this section, we present our strategy to identify the necessary statistics to evaluate the theoretical framework developed in Section 1.2. To this end, we will first present the electricity pricing mechanism used in Zurich and then show how we exploit the price variation to identify the statistics outlined above.

9. Source: Department of Statistics, Canton Zurich: https://statistik.zh.ch, last visited: December 20, 2021 and Swiss Federal Office of Statistics: https://www.bfs.admin.ch, last visited: December 20, 2021

1.4.1 The Pricing Mechanism

In Switzerland, like in many other countries, households do not pay the same price for electricity during each hour of the day. At night (during the night tariff *NT*), electricity consumption is substantially cheaper than during the day (day tariff, or *DT*) (during our sample period, this price difference was 0.0973 CHF/kWh, with households paying on average approximately 0.21 CHF per kWh during the DT and 0.11 CHF per kWh during the NT). This has efficiency reasons: Electricity cannot easily be stored: It has to be produced when it is actually needed. Electricity plants and grids thus have to be prepared for peak hours, i.e. the time, when electricity demand is highest. The peak hours are a substantial cost driver for operators, both on the electricity production and on the transportation side, as plants are usually not flexible enough to quickly adapt their production to demand. In order to smooth aggregate electricity consumption over the day (and thus capacities of the power grids) by incentivising electricity consumption during off-peak hours, plants set lower prices when aggregate demand is low (during the night, i.e. 10 p.m. - 6 a.m.). The price differential between the DT and NT is thus foremost an incentive for households to shift electricity consumption towards the NT.
The DT lasts from Monday to Saturday, from 6 a.m. to 10 p.m. - leaving Sundays in the NT and making electricity consumption during the day on Sundays substantially cheaper than on Saturdays. The reason for this is the fact that heavy industry rests on Sundays, as opposed to Monday through Saturday. This generates an excess supply of energy, so that plants lower electricity prices on Sundays as well. This provides us with the exogenous price variation we need for our theoretical framework: While DT-consumption on Saturday costs between 0.1945 and 0.243 CHF/kWh (depending on the electricity product[10]), it costs 0.0973 CHF/kWh less on Sundays, while the nighttime price stays constant across the two days. Figure 1.3 illustrates the price variation used in our framework. On average, DT-consumption costs roughly 0.21 CHF per kWh for each day of the week (in contrast, NT-consumption only costs around 0.11 CHF per kWh on average), whereas it costs only around 0.11 CHF per kWh on Sundays.

Note that we will focus our analysis on the price variation between Saturday and Sunday, as we believe these days are similar enough to allow for a direct comparison.[11] In order to identify the salience parameter from equation (1.11) we only need to assume that any differences in electricity consumption between Saturday and Sunday are structural and do not differ based on treatment status (except, of course, any load shifting that is done to exploit the pricing mechanism). This assumption is fairly weak and easy to justify: In a difference-in-differences setup it

10. For a detailed overview of the different tariffs, see Table 1.A.7 in Appendix 1.A.
11. However, we also present evidence on the effect of price variation between Sunday and Monday using an empirical approach that is adapted to account for the structural differences between Sunday and Monday (as opposed to Saturday-Sunday price variation).

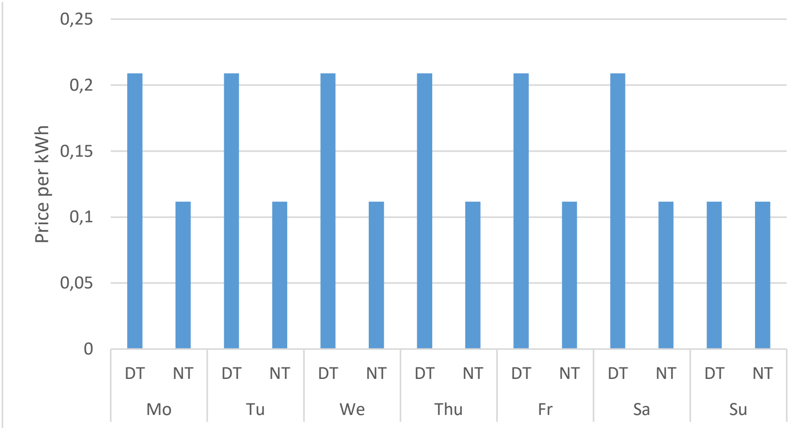

Figure 1.3. Price Chart over Days of the Week

This figure presents the average pricing schedule for households in CHF per kWh. The price of electricity is calculated as a weighted average of the prices of the oekopower, naturpower, and mixpower-electricity products. Daytime electricity consumption is more expensive than nighttime electricity throughout the week and on Saturdays. On Sundays, however, the daytime electricity price stays at the nighttime-level. The non-linearity generates a price difference in the DT-electricity price between Saturday and Sunday, while the NT-electricity price stays constant over both days.

would be the equivalent to the parallel-trends assumption.

Identifying the salience parameter from equation (1.8), however, requires that differences in electricity consumption between Saturday and Sunday are only due to the price differences between the two days. Note that this assumption is still justifiable, but much stronger. It would be unlikely that it would hold it we compared a weekday to Sunday: Households behave differently during weekdays, as people work or are out of the house for other reasons. However, we can assume that energy consumption behavior on Saturdays and Sundays allows for a direct comparison. Figure 1.4 supports our line of though: Here, we plot the average pre-treatment electricity load profiles over the day for different days of the week. We observe substantially different load profiles over the day for the weekend and weekdays: While on weekdays, electricity consumption is lower for almost every time of the day, it jumps up between 5 and 7 p.m. as people are preparing to leave for work. After this initial early jump, electricity consumption stagnates until 12 p.m., whereas on weekends, it continuously increases. The load profiles for Saturdays and Sundays, however, look alike. They indicate that people get up later during the weekend, but stay at home, resulting in higher average consumption over the day. Their profiles look very similar in shape, though consumption is higher in the evening on Sundays than on Saturdays. We argue that the differences between the Saturday and Sunday load profiles are mostly due to the existing price differences between the two days. In all three graphs, we observe a sharp increase after 4 p.m. until 8 p.m. as people return from work, start to prepare dinner etc. The higher level in consumption on weekends, which is especially pronounced between 8 a.m. and 4 p.m. can be

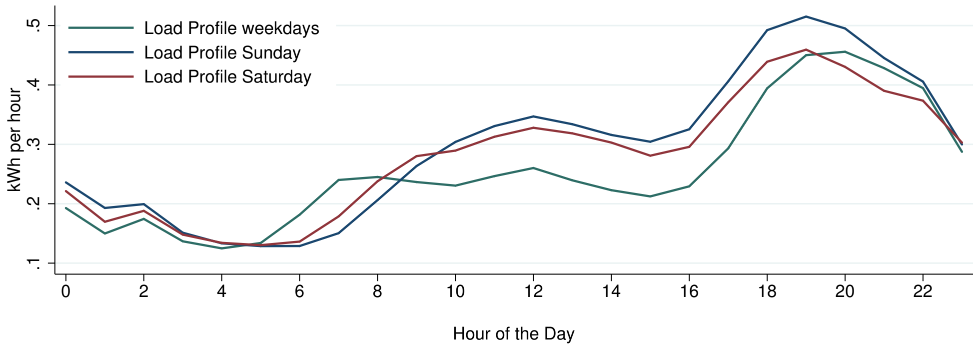

Figure 1.4. Electricity Load Profiles for Different Days of the Week

This figure presents the average electricity load profiles in kWh per hour on Saturday, Sunday and weekdays before IHD-installation. For this graph, we aggregated the consumption data of households on an hourly level. Note that the unit of measurement is kWh/hour = kW for each hour of the day.

explained by more people being at home instead of being at work (note that this pattern, though less pronounced, even extends up to 6 or 8 p.m.). Hence, while weekends are certainly not representative of a household's energy consumption profile in terms of consumption, we have to focus on Saturdays and Sundays to make sure we do not mix up price- and weekday-effects.[12]

1.4.2 Graphical Evidence

In this subsection, we will present graphical evidence that will illustrate the conservation efforts and load shifting households undergo after being introduced to the In-Home-Displays. This will also pave the path to our identification strategy for the salience parameter θ, on which we will elaborate in Section 1.4.5.

Figure 1.5 illustrates the treatment effect of installing an IHD using the average daily load profiles for households before and after IHD-installation (counting the first 50 days after IHD-installation). Note that we can clearly see a reduction in electricity consumption for Saturdays, especially during the DT from 6 a.m. to 10 p.m., whereas the effect of the IHDs on Sunday consumption seems to be somewhat ambiguous. The differences in consumption before and after IHD-installation may seem small at first, but remember that the unit of measurement is kWh/hour = kW for each hour of the day. In order to take a closer look at the treatment and price

12. Additionally, we will exclude households with *Solartop*-pricing-rate, as these rates have a much higher price per kWh and do not discriminate between DT- and NT-consumption. Also note that the observations excluded by leaving out Solartop-households only make up 0.37% of the total sample.

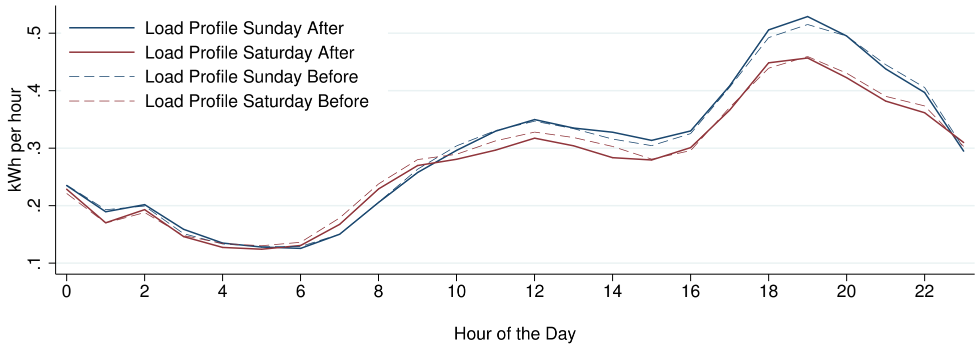

Figure 1.5. Change in Electricity Load Profiles After IHD-installation

The figure presents the average electricity load profiles before and after IHD-installation (counting the first 50 days after IHD-installation) for Saturdays and Sundays. Note that the unit of measurement is kWh/hour = kW for each hour of the day.

effects in the data, we will now zoom in on the deviations of the IHD-households from the control households (that is, households with an installed smart meter, but no feedback through an IHD, yet). We normalize all observations by the average per-hour-consumption of a household to filter out individual fixed effects and by the average per-hour-consumption in the control group to filter out time fixed effects.[13] Remember that *control group* refers to households with smart meter, but without IHD at the corresponding date and hour of the day. Figure 1.6 illustrates the deviations from the hourly mean in the first 50 days after IHD-installation.

The two dashed lines denote the normalized average consumption of the control group, which, unsurprisingly, fluctuates around 0, as this is the group we used to normalize the data.[14] Looking at the two other graphs, we first see massive reductions in energy consumption, which, in light of Figure 1.4 are most pronounced during the peak consumption hours (especially between 6 p.m. and 9 p.m.). We also observe a slight increase in consumption between 12 a.m. and 4 a.m. or 8 a.m., respectively for Saturday and Sunday. This suggests electricity load shifting from DT- to NT-hours: Following IHD-introduction, households exploit the energy pricing mechanism by shifting electricity consumption from the expensive DT to the cheaper NT. Although Figure 1.4 clearly shows higher baseline consumption on Sundays compared to Sat-

13. Note that including fixed-effects makes our results more precise, but can slightly alter the deviation-graphs due to the unbalanced IHD-roll-out. Thus, while the deviation-graphs indeed represent the results found in the load profile-graph in Figure 1.5, they may differ in small details.

14. Note that the "before" graphs do not exactly equal 0 because we filtered out fixed effects before restricting our sample to the 50 days before and after IHD-installation.

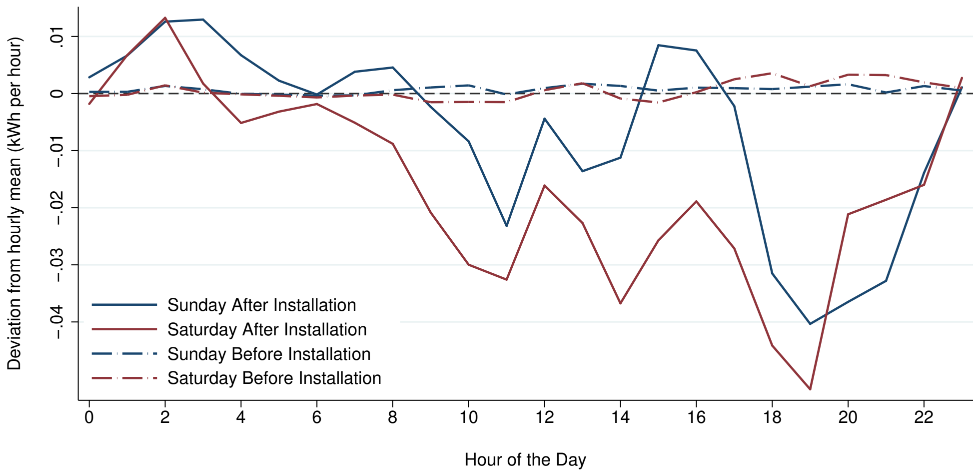

Figure 1.6. Deviations from the Hourly Average in kWh per Hour

Deviations from the hourly mean in the first 50 days after IHD-installation on Saturdays and Sundays. The two dashed lines denote the normalized average consumption of the control group, that is, households with smart meter, but no IHD, yet.

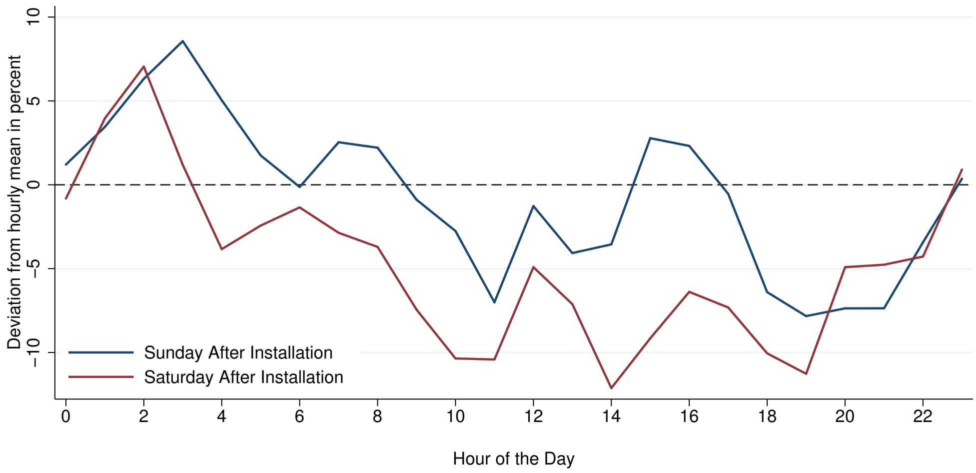

Figure 1.7. Deviations from the Hourly Average of Untreated Households in Percent of Pre-Treatment Energy Consumption

This figure presents the deviations from hourly average consumption of pre-IHD-installation households in percent for Saturdays and Sundays during the first 50 days after IHD installation. We can see that during the daytime tariff, a nearly constant percentage of hourly consumption is conserved after IHD-installation. This explains the larger absolute decrease in consumption during the evening hours following IHD-installation.

urdays, we observe the largest reductions in energy consumption on Saturdays, not on Sundays.[15]

15. In Appendix 1.C, we additionally illustrate the average treatment effect for each hour of the day using a Difference-in-Difference specification on Saturdays, Sundays and weekdays, respectively.

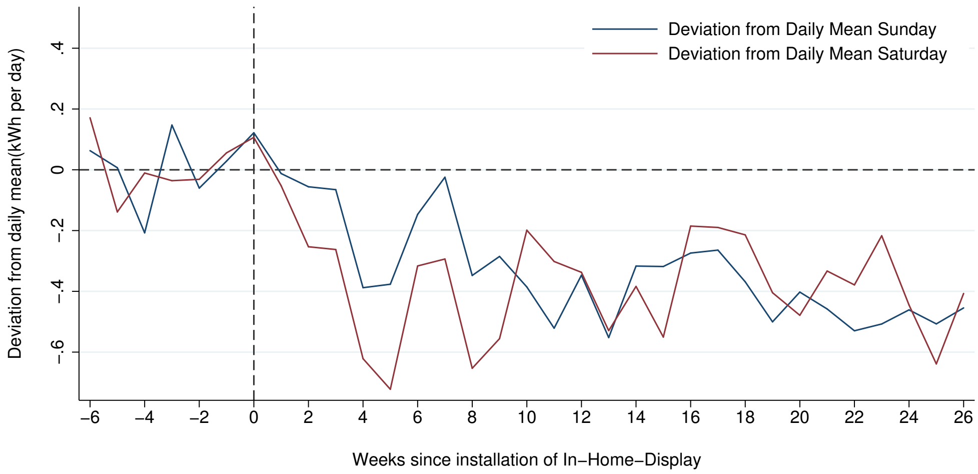

Figure 1.8. Deviations from the Daily Average in kWh per Day

This figure presents the deviations from the pre-IHD-installation daily average in kWh per day and illustrates the persistence of the treatment effect of IHD-installation.

Figure 1.7 reveals the reason for this: At each hour during the DT, households reduce hourly electricity consumption by a constant fraction of roughly 5 to 10 percent on Saturdays and considerably less on Sundays (though still substantial). This is first evidence for a stronger reaction to feedback if energy prices deliver an additional incentive. We would expect higher baseline consumption on Sundays to result in higher reductions on Sundays, given that higher consumption leaves more scope and larger incentives for conservation efforts, but the highest reductions are achieved during the expensive Saturday-DT.

Figure 1.8 illustrates the persistence of the treatment effect for Saturdays and Sundays, respectively. Again, the results (which include household-fixed effects) are normalized by the daily mean for each date. After IHD-installation, we see an immediate drop of electricity consumption, which persists even beyond the 50 days ($\approx$ 7 weeks), to which we restricted our previous graphs. We also observe that the response for Saturdays seems to be a little more pronounced than the response for Sundays, though the difference is not significant.

1.4.3 Putting Framework and Data Together

Looking at Section 1.2 and the respective equations to identify the salience parameter θ, there are several statistics we have to recover from the data:

- Δx_1^*: The average effect on DT-consumption of correcting the salience bias. Using our assumption that installing an IHD fully corrects the salience bias, this becomes the average treatment effect of IHD-installation on DT-consumption.

- ψ_{11}^{FB}: The reaction of DT-consumption (under feedback) to a change in the daytime price. In our framework, this is constant in x_1 (and x_2), such that $\psi_{11}^{FB} = \frac{\Delta x_1^{FB,p_1}}{\Delta p_1}$. We thus only need to identify $\Delta x_1^{FB,p_1}$, i.e. the increase in DT-consumption between Saturday and Sunday due to the decrease in p_1.
- ψ_{11}^{*}: The reaction of DT-consumption *without* feedback to a change in the daytime price. Like above, we only need to identify $\Delta x_1^{*,p_1}$
- ψ_{21}^{FB}: The reaction of NT-consumption (under feedback) to a change in the daytime price. This will give us an idea of the relationship of DT- and NT-consumption (i.e. if and to what degree they are substitutes or complements). As above, our framework implies that the derivative is constant in x_1 and x_2, such that $\psi_{21}^{FB} = \frac{\Delta x_2^{FB}}{\Delta p_1}$. Again, we thus only need to identify Δx_2^{FB}, i.e. the reduction/increase in NT-consumption between Saturday and Sunday due to the decrease in p_1.

1.4.4 Reduced Form Evidence

Before turning to a General Method of Moments (GMM) framework in order to estimate our model, we will first present the reduced-form results of a simple difference-in-differences estimation. In order to avoid implementation and coordination problems, the smart meters and IHDs had to be rolled out over the course of the year, meaning that each household belonged to a squadron of households receiving the IHD at a certain time of the study (for most households, this was approximately eight weeks after the installation of the smart meter). This staggered roll-out has some important implications for our empirical design: In order to identify the necessary statistics, especially Δx_1^{*}, i.e. the average treatment effect of installing an IHD on DT-consumption, we cannot simply compare average consumption before and after IHD-installation, as this would mix up the actual treatment effect with time effects on electricity consumption. For example, higher use of air conditioning during the summer or leaving the lights on during the winter would give us different treatment results for the same household, depending on the date of IHD-installation. Hence, we need to filter out the time effects induced by the staggered roll-out. For this, we need a control group, consisting of untreated households (i.e. who do not receive feedback yet). As noted in Section 1.3.1, the *control group* for our empirical strategy will consist of households with an installed smart meter, but no installed IHD. The IHD roll-out was distributed across squadrons over the period from January 2011 to July 2012 (according to the timeline from Figure 1.1). The point in time T on the x-axis does not refer to a single date that is identical for all households but rather a squadron-specific date of IHD-installation. These dates (and hence the weeks of baseline consumption recording) are distributed over the course of the study, allowing us to include time-fixed effects in our analysis. Note that assignment to the IHD-installation and hence treatment timing was random. Figure 1.9 illustrates the roll out. We plot the share of households with smart meter, but without IHD. Note that

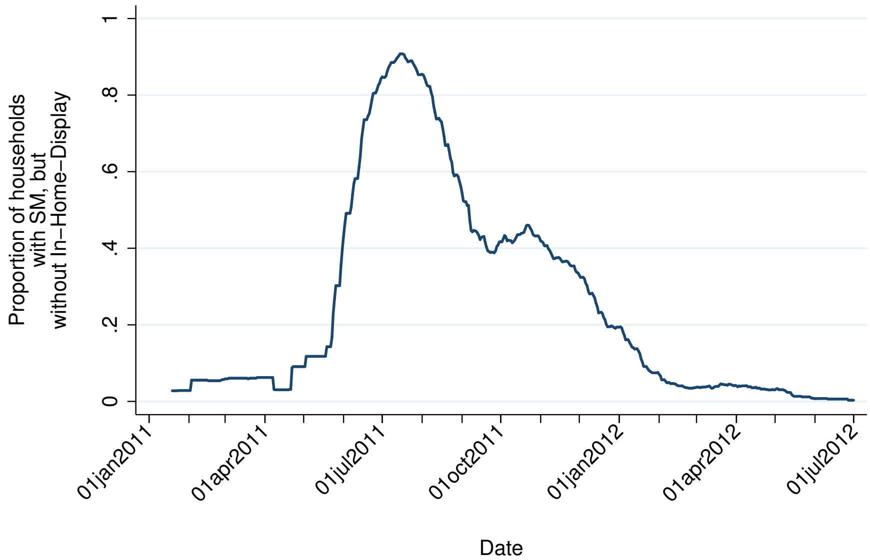

Figure 1.9. Share of Households with Smart Meter, but without IHD

This graph displays the share of households with an installed smart meter, but no IHD installed, yet.

the share does not, as one might expect, fall monotonically, but first increases due to smart meter installation and then decreases due to IHD-installation. The graph thus illustrates the size of our control group.

Hence, in order to identify the average treatment effect of IHD-introduction (and thus, resolving the salience bias) for each combination of day-of-the-week and time-of-the-day, we conduct a simple difference-in-differences regression on each corresponding subset of observations. Note that we will subsume the weekdays from Monday to Friday as "Weekdays" instead of analyzing each day separately. Before linking our data to the theoretical framework, we run the following difference-in-differences specification for Saturdays, Sundays and Weekdays, respectively, and separately for daytime (DT) and nighttime energy consumption (NT).

$$y_{it}^{Sa,DT} = \alpha_i^{Sa,DT} + \beta_{w(t)}^{Sa,DT} + \delta^{Sa,DT} D_{i,t} + \epsilon_{it}^{Sa,DT}$$

That is, if we restrict our sample to observations on Saturday during the DT, the coefficient on the dummy $D_{i,t}$ (which is 1 if observation t of individual i is recorded *after* IHD introduction) gives us the average treatment effect of IHD-introduction on Saturdays during the DT. The control group in this framework consists of households with an installed smart meter, but no feedback through IHDs. We include $\alpha_i^{Sa,DT}$ and $\beta_{w(t)}^{Sa,DT}$ for household- and time fixed effects (on a weekly level). We can conduct the regression on each set of observations corresponding to each of the six combinations between Saturday, Sunday and Weekdays and Daytime and Nighttime tariff. This gives us the following set of regression equations:

$$y_{it}^{Sa,DT} = \alpha_i^{Sa,DT} + \beta_{w(t)}^{Sa,DT} + \delta^{Sa,DT} D_{i,t} + \epsilon_{it}^{Sa,DT} \tag{1.16}$$
$$y_{it}^{Sa,NT} = \alpha_i^{Sa,NT} + \beta_{w(t)}^{Sa,NT} + \delta^{Sa,NT} D_{i,t} + \epsilon_{it}^{Sa,NT} \tag{1.17}$$
$$y_{it}^{Su,DT} = \alpha_i^{Su,DT} + \beta_{w(t)}^{Su,DT} + \delta^{Su,DT} D_{i,t} + \epsilon_{it}^{Su,DT} \tag{1.18}$$
$$y_{it}^{Su,NT} = \alpha_i^{Su,NT} + \beta_{w(t)}^{Su,NT} + \delta^{Su,NT} D_{i,t} + \epsilon_{it}^{Su,NT} \tag{1.19}$$
$$y_{it}^{We,DT} = \alpha_i^{We,DT} + \beta_{w(t)}^{We,DT} + \delta^{We,DT} D_{i,t} + \epsilon_{it}^{We,DT} \tag{1.20}$$
$$y_{it}^{We,NT} = \alpha_i^{We,NT} + \beta_{w(t)}^{We,NT} + \delta^{We,NT} D_{i,t} + \epsilon_{it}^{We,NT} \tag{1.21}$$

Then, the δ-coefficients deliver the average treatment effect of IHD-introduction on each day-of-the-week and time-of-the-day. Table 1.2 presents the results from the regressions described above.

As in Figure 1.6, we observe large, highly significant reductions in average energy consumption during the DT for all three day categories following IHD-introduction: The strongest conservation potentials are achieved on Saturdays, when baseline consumption is high compared to weekdays, and the energy pricing structure provides additional incentives for conservation efforts. For Saturdays during the DT, consumption decreased by 0.316 kWhs ($\approx$ 6 percent) on average. Although baseline consumption is lowest during weekdays, we observe the second largest reduction in consumption on these days during the DT ($\approx$ 0.247 kWhs, or 5% on average). Note that the high daytime price provides additional incentives for energy consumption reduction despite the smaller scope for conservation efforts. The smallest, though still significant, reduction in energy consumption is achieved on Sundays. On Sundays, we estimate a reduction in DT consumption of 0.203 kWhs ($\approx$ 4 percent) on average, even though baseline consumption is highest and hence, the scope for energy conservation should be the larger on Sundays than on Saturdays and weekdays. However, as DT consumption is substantially cheaper on Sundays, households are faced with weaker financial incentives for conservation efforts, likely causing the weaker treatment effect of IHD-introduction.
The results are especially interesting in the light of Figure 1.7, where we see that following feedback provision, households reduce energy consumption by a nearly constant fraction during the DT, such that, if price incentives were not a large driver of our results, we would expect the treatment effect to be monotonous in baseline electricity consumption with respect to Saturdays, Sundays and weekdays.
We do not find any evidence of increased nighttime electricity consumption after IHD-introduction for any day of the week, so we find no reduced form evidence

Table 1.2. Average Treatment Effect for Each Day of the Week and Time of the Day

	(1)	(2)	(3)	(4)	(5)	(6)
Dependent Variable	Consumption	Consumption	Consumption	Consumption	Consumption	Consumption
Treatment Effect Saturday DT	-0.316***					
	(0.072)					
Treatment Effect Saturday NT		-0.042				
		(0.025)				
Treatment Effect Sunday DT			-0.203***			
			(0.069)			
Treatment Effect Sunday NT				-0.048*		
				(0.025)		
Treatment Effect Weekday DT					-0.247***	
					(0.048)	
Treatment Effect Weekday NT						0.007
						(0.020)
Individual-FEs	Yes	Yes	Yes	Yes	Yes	Yes
Week-FEs	Yes	Yes	Yes	Yes	Yes	Yes
Observations	43'554	43'590	43'420	43'474	217'915	218'137
R^2_{within}	0.055	0.010	0.059	0.009	0.045	0.007
$R^2_{between}$	0.012	0.000	0.015	0.002	0.014	0.000
$R^2_{overall}$	0.029	0.003	0.030	0.001	0.024	0.003
Clusters	978	978	977	977	979	979

Notes: This table presents results from the DiD-regressions based on equations (1.16)-(1.21). Dependent variable: Daily DT or NT electricity consumption, respectively. The estimated coefficients denote mean deviations from pre-treatment baseline consumption, based on indicators for each combination of day-of-the-week and DT/NT-tariff. The control group in this framework consists of households with an installed smart meter, but no feedback through IHDs. Each column reports the mean deviation (= average treatment effect) for each day-of-the-week based on equations (1.16) - (1.21). Each treatment effect was estimated in a separate DiD regression. The control group in each specification consists of households with an installed smart meter, but which do not receive feedback through an IHD, yet. All specification include week-of-sample- and household fixed effects. Standard errors (in parentheses) are clustered at the household level.
*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

of load shifting using our difference-in-differences specifications. However, we do not observe any significant reductions in energy consumption during the night on weekdays and Saturdays, either. More interestingly, we find significant evidence of reductions in nighttime electricity consumption of 0.048 kWhs during the NT on Sundays (when the pricing schedule provides no price differential between DT- and NT consumption), but not on Saturdays and weekdays.
Of course, looking at Figure 1.6, we would expect at least some evidence of load shifting. However, note that our reduced form analysis is conducted on the daytime-nighttime-level rather than analysing each hour of the day separately. Hence, it may be subject to spillover effects from one hour to another. We will further explore this channel in Section 1.4.6 and in Appendix 1.B by excluding the cutoff-hour between 9:30 and 10:30 p.m.

As the overidentification of the salience parameter θ in our framework allows for θ to be identified via a GMM-Setup, we will now turn to a GMM identification strategy to evaluate the model set up in Section 1.2.

1.4.5 GMM Estimation

Our theoretical framework from Section 1.2 yields several specifications of the derivatives of electricity consumption with respect to the electricity prices and the salience parameter θ. As mentioned in Section 1.2.3, we can identify the salience parameter θ via two independent methods. Since our model is overidentified, we can thus invoke a GMM-estimation framework.
The framework from Section 1.2 gives us

$$\frac{\partial x_1^*}{\partial \theta} = \left(p_1 \frac{\partial x_1^{FB}}{\partial p_1} + p_2 \frac{\partial x_2^{FB}}{\partial p_1} \right) \tag{1.22}$$

$$\frac{\partial x_2^*}{\partial \theta} = \left(p_2 \frac{\partial x_2^{FB}}{\partial p_2} + p_1 \frac{\partial x_2^{FB}}{\partial p_1} \right) \tag{1.23}$$

and

$$\frac{\partial x_1^*}{\partial p_1} = \theta \frac{\partial x_1^{FB}}{\partial p_1} \tag{1.24}$$

$$\frac{\partial x_2^*}{\partial p_1} = \theta \frac{\partial x_2^{FB}}{\partial p_1} \tag{1.25}$$

While our data is generated as follows:

$$\begin{aligned} x_{1it} \; = \; & \alpha_i + \delta_{w(t)} + x_{Sa,DT,before} + \frac{\partial x_1^*}{\partial \theta}(1-\theta) \cdot \mathbf{1}_{Sa,DT,After} \\ & + \frac{\partial x_1^*}{\partial p_1} \Delta p_1 \cdot \mathbf{1}_{Su,DT,Before} \\ & + \left[\frac{\partial x_1^*}{\partial \theta}(1-\theta) + \frac{\partial x_1^{FB}}{\partial p_1} \Delta p_1 \right] \cdot \mathbf{1}_{Su,DT,After} \\ & + \epsilon_{1it} \end{aligned} \tag{1.26}$$

and

$$\begin{aligned} x_{2it} \; = \; & \alpha_i + \delta_{w(t)} + x_{Sa,NT,before} + \frac{\partial x_2^*}{\partial \theta}(1-\theta) \cdot \mathbf{1}_{Sa,NT,After} \\ & + \frac{\partial x_2^*}{\partial p_1} \Delta p_1 \cdot \mathbf{1}_{Su,NT,Before} \\ & + \left[\frac{\partial x_2^*}{\partial \theta}(1-\theta) + \frac{\partial x_2^{FB}}{\partial p_1} \Delta p_1 \right] \cdot \mathbf{1}_{Su,NT,After} \\ & + \epsilon_{2it} \end{aligned} \tag{1.27}$$

The data generating process is illustrated in Figure 1.10. Excluding household- and time fixed effects α and δ and restricting the sample to observations on Saturdays and Sundays, the difference in average energy consumption on Saturdays during the DT *before* IHD-installation and *after* IHD-installation (i.e. the average treatment effect of IHD-introduction) is denoted by Δx_1^*. The average difference in energy consumption between *Saturday* and *Sunday* during the DT *before* IHD-installation, i.e. the average pre-treatment price effect from lowering the DT-energy price from 0.21 CHF per hour to 0.11 CHF per hour, is denoted by $\Delta x_1^{*,p_1}$. Similarly, the average difference in energy consumption between *Saturday* and *Sunday* during the DT *after* IHD-installation is denoted by $\Delta x_1^{FB,p_1}$. The data generating process for NT-energy consumption works analogously.

The rationale behind equations (1.26) and (1.27) is as follows: Household i's electricity consumption on date t during the day x_{1it} and during the night x_{2it} is determined by the individual household's baseline consumption α_i, time-fixed effects $\delta_{w(t)}$, an indicator for the day of the week and an indicator for IHD-installation status (*before* or *after*) along with the average difference in consumption associated with each day-of-the-week-tariff combination. For each time of the day $\{DT, NT\}$, there are four possible combinations of day-of-the-week and treatment status, each one with their own indicator.[16]

16. Note that for this analysis, we restrict the sample to observations on Saturdays and Sundays to only exploit the price variation between Saturdays and Sundays.

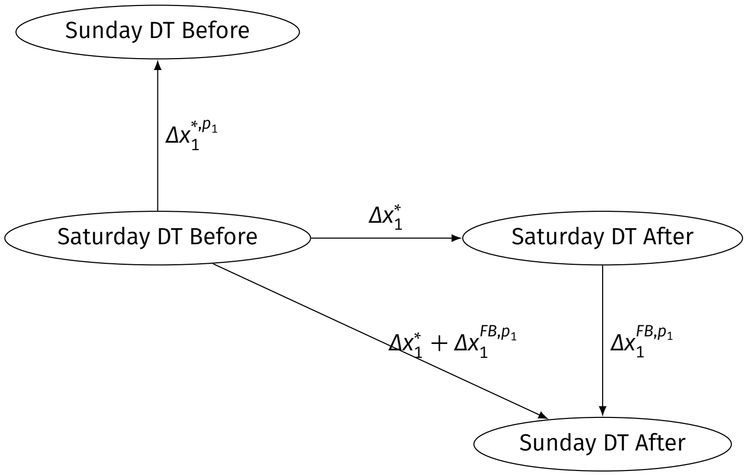

Figure 1.10. The Data Generating Process

Diagram of the data generating process. For example, the difference in electricity consumption between Saturday during the DT *before* intervention and Saturday during the DT *after* intervention (i.e. the DT treatment effect) is denoted by Δx_1^*. Similarly, the difference between consumption on *Saturday* during the DT before intervention and *Sunday* during the DT before intervention (i.e. the demand response to the price change between Saturday and Sunday during the DT *before* intervention) is denoted by $\Delta x_1^{*,p_1}$. By the same logic, the same demand response *after* intervention is denoted by $\Delta x_1^{FB,p_1}$

Assuming constant derivatives and plugging in equations (1.22), (1.23), (1.24) and (1.25) gives us

$$\begin{aligned}
x_{1it} = {} & \alpha_i + \delta_{w(t)} + x_{Sa,DT,before} + \left(p_1 \frac{\partial x_1^{FB}}{\partial p_1} + p_2 \frac{\partial x_2^{FB}}{\partial p_1} \right)(1-\theta) \cdot \mathbf{1}_{Sa,DT,After} \\
& + \frac{\partial x_1^*}{\partial p_1} \Delta p_1 \cdot \mathbf{1}_{Su,DT,Before} \\
& + \left[\left(p_1 \frac{\partial x_1^{FB}}{\partial p_1} + p_2 \frac{\partial x_2^{FB}}{\partial p_1} \right)(1-\theta) + \frac{\partial x_1^{FB}}{\partial p_1} \Delta p_1 \right] \cdot \mathbf{1}_{Su,DT,After} \\
& + \epsilon_{1it}
\end{aligned} \tag{1.28}$$

and

$$\begin{aligned}
x_{2it} = {} & \alpha_i + \delta_{w(t)} + x_{Sa,NT,before} + \left(p_2 \frac{\partial x_2^{FB}}{\partial p_2} + p_1 \frac{\partial x_2^{FB}}{\partial p_1} \right)(1-\theta) \cdot \mathbf{1}_{Sa,NT,After} \\
& + \frac{\partial x_2^*}{\partial p_1} \Delta p_1 \cdot \mathbf{1}_{Su,NT,Before} \\
& + \left[\left(p_2 \frac{\partial x_2^{FB}}{\partial p_2} + p_1 \frac{\partial x_2^{FB}}{\partial p_1} \right)(1-\theta) + \frac{\partial x_2^{FB}}{\partial p_1} \Delta p_1 \right] \cdot \mathbf{1}_{Su,NT,After} \\
& + \epsilon_{1it}
\end{aligned} \tag{1.29}$$

Where $x_{Sa,DT,before}$ denotes energy consumption on Saturdays before IHD-installation during the DT and $x_{Sa,NT,before}$ denotes energy consumption on Saturdays before IHD-installation during the NT, as we left the corresponding indicators out of the equation as a baseline category. That is, equation (1.28) denotes the moment condition for DT consumption and equation (1.29) denotes the moment condition for NT consumption. We then fit equations (1.28) and (1.29) in a GMM estimation framework. Finally, we cluster all standard errors at the household level. Additionally, we conduct a difference-in-differences analysis to identify the statistics needed to calculate the salience parameter θ. The framework and results are presented in the Appendix 1.F. We use the results from the DiD-analysis as starting values for the GMM analysis in the following section.

1.4.6 Main Results

The results from our GMM analysis are presented in Table 1.3. We estimate the salience parameter θ to be 0.681. That is, in our sample, we estimate that households behave as if they only perceive 68 percent of their electricity costs. The estimated derivative of DT-energy consumption with respect to its price is −4.370. That is, increasing the DT-electricity price by 1 CHF decreases DT-energy consumption by 4.37 kWhs per day on average. Note that we assume this derivative to be constant in energy consumption. In our sample, average post-treatment electricity consumption during the DT on Saturdays is 4.60 kWhs. We can thus calculate a post-treatment price elasticity of −0.200.
Our estimate of -0.703 for the derivative of NT-energy consumption with respect to its own price is not statistically significant. There are two simple explanations for the ambiguous estimate of this parameter: First, the data does not allow for it to be observed directly. The point estimate is a “byproduct” of our estimation framework, but note that we do not observe any variation in p_2. The second explanation lies in the baseline of NT-energy consumption: With a much lower baseline consumption during the NT, households have little scope for consumption adjustment during the night.

We estimate a cross-price derivative of NT consumption with respect to the DT price of -0.660. The sign of the estimated derivative may appear puzzling: Basic economics tells us that a negative sign of the cross-derivative indicates complements, whereas we would expect DT- and NT-energy consumption to be substitutes. However, as Figures 1.5 and 1.6 show, most energy consumption adjustment (following IHD-introduction or price variation) occurs during high consumption hours between 5 and 11 p.m. This also means that some adjustments that address DT-consumption are partially attributed to nighttime consumption (e.g. using the washing machine

Table 1.3. GMM Results

	Estimate
θ	0.681***
	(0.072)
ψ_{11}^{FB}	-4.370***
	(0.442)
ψ_{21}^{FB}	-0.660***
	(0.107)
ψ_{22}^{FB}	-0.703
	(0.659)
Test for $\theta = 1$:	0.000
Hansen's J χ^2	2.718
p-value Hansen's J:	0.257
Observations	163'028
Clusters	907

Notes: The table presents the results for the salience parameter θ and the price sensitivities ψ from the GMM estimation based on the moment conditions that underlie the GMM framework, 1.28 and 1.29. Standard errors (in parentheses) are clustered at the household level. To estimate θ, a weighted average of different prices (depending on the electricity product) was used to approximate p_1 and p_2. A Wald-test tests whether the estimate for θ is significantly different from 1. Hansen's J statistic is used to test the validity of the overidentifying restrictions in a GMM model. We do not reject the null hypothesis that the model is correctly specified.
*** $p < 0.01$ ** $p < 0.05$ * $p < 0.1$

or dishwasher during a different time of the day). In other words: During a short time window, DT- and NT-energy consumption can be seen as complements. To illustrate this, consider a dishwasher that usually runs from 9:45 to 10:30 p.m. (e.g. after people had dinner, watched a movie and start to prepare to go to bed). Then, between 9:45 and 10:30 p.m., DT and NT electricity consumption for the dishwasher are indeed complements, as it has to run the full time to clean the plates. Similar mechanisms hold, for example, for washing machines. If an energy intensive device runs during the DT-cutoff time at 10 p.m. (which is likely given the high average consumption during this time), this would render DT- and NT-energy consumption complements rather than substitutes.[17] Any reductions in DT-consumption around the cutoff would then spill over into the night tariff, where their relative magnitude would blur the "actual" reduction effects of the NT due to the higher consumption baseline at 10 p.m. compared to the rest of the NT. In Appendix 1.B, we thus run

17. Our theoretical framework can easily be extended to a more realistic three-good case, with the third good being electricity consumption in the evening. Of course, since we do not observe price variation for certain hours of the day, we would not be able to apply this framework to our real setting.

our regressions again, leaving out all observations between 9:30 and 10:30 p.m. The results support our intuition, as we observe an estimate for ψ_{21}^{FB} of a smaller magnitude, which is still significantly negative. There remains some uncertainty regarding the mechanisms driving the substitution effects of DT- and NT-energy consumption. Note, however, that our theoretical framework allows to identify θ via two methods: As the ratio of the price sensitivities before and after feedback and as a complex function of the IHD-treatment effect and the price elasticities of DT- and NT-consumption.[18]. While the first method does not require any knowledge of the cross-derivative of NT energy consumption with respect to the DT energy price at all, method 2 indeed uses ψ_{21}^{FB}. However, even here, the effect of ψ_{21}^{FB} barely carries any weight, relative to Δx_1^* and ψ_{11}^{FB} as it is scaled down by the low nighttime price. Hence, even if our estimate for ψ_{21}^{FB} was flawed, it would barely affect our estimate for θ. Lastly, Hansen's J-statistic (which is used to determine the validity of the overidentifying restrictions in a GMM model) is not significant, indicating that our model is well-specified.

1.4.7 Exploring Heterogeneity in the Results

A natural question that comes to mind when looking at our results is how they change for different households. Our data contains information on household characteristics such as household size, income, tenancy status, environmental attitudes and more. Furthermore, the IHDs record the number of interactions with the displays and hence, provide a measure of IHD-usage. Table 1.4 reports the dimensions along which we explore heterogeneity in our results and provides the correlation structure between the heterogeneity variables.

An important implication of our framework might appear puzzling: consumption derivatives with respect to energy prices and salience are implied to be constant in current energy consumption. However, a household with a baseline consumption of 10 kWhs per day should clearly exhibit a higher conservation potential than a household with a baseline consumption of 4 kWhs per day due to the larger scope (and incentives) for conservation efforts. Note that without our assumption, we would not be able to identify the key parameters in our model. But in order to explore the potential shortcomings of our framework, we conduct our analysis on different subsets of the sample, based on the terciles of the initial (baseline) consumption distribution of all households in the smart-metering treatment group, to see how the estimates differ. In order to avoid endogeneity concerns regarding the baseline split, we conduct the baseline consumption tercile split for this analysis based on the days we left out of our main analysis, i.e. Monday through Friday.

18. See equations (1.8) and (1.11)

Table 1.4. Correlation Structure Between Key Characteristics

	IHD-Interactions	Baseline Consumption	Income	Household Size	Education
IHD-Interactions	1				
Baseline Consumption	0.0872**	1			
Income	0.0396	0.236***	1		
Household Size	0.0585	0.453***	0.303***	1	
Education	-0.0238	-0.128***	0.243***	0.0653*	1

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Notes: The table presents unconditional correlations between several household characteristics. *Income* levels correspond to the levels depicted in Figure 1.A.4 (< 3000*CHFs/Month*, $3000 - 3999$*CHFs/Month*, $4000 - 5999$*CHFs/Month*, $6000 - 6999$*CHFs/Month*, $7000 - 7999$*CHFs/Month*, $8000 - 8999$*CHFs/Month*, $9000 - 9999$*CHFs/Month*, $10000 - 11999$*CHFs/Month*, $12000 - 14999$*CHFs/Month*, $15000 +$ *CHFs/Month*), corresponding to the values 1-11. *Education* levels are ordered according to the order pictured in Figure 1.A.6 (*No Education, Obligatory Schooling, 2-3 Years General Schooling, Apprenticeship, Federal Vocational Baccalaureat, A-Level, College, University and Others*), excluding the category "other", corresponding to the values 1-8.
* $p < 0.1$, ** $p < 0.05$, ***$p < 0.01$

Additionally, we take a look at households that frequently used the IHD compared to households which used them less frequently. In light of the concerns expressed by Fischer (2008) and Buchanan, Russo, and Anderson (2015) regarding the importance of user interaction, we expect the results to be amplified for high IHD usage households, but not rendered insignificant for low IHD usage households.

Table 1.5 reports the results from our heterogeneity analysis. Column (1) reports the results from our baseline sample as in Table 1.3. Columns (2), (3) and (4) show the results from the baseline consumption heterogeneity analysis. For this we split up the sample into households with an average daily energy consumption above and below the 33 and 66 percentile. Figure 1.2 illustrates the tercile split.
As we would expect given the large scope for conservation efforts, high baseline-consumption households consumption tend to be the most price sensitive after IHD-introduction. The pattern we observe in our heterogeneity analysis suggests that post-feedback price sensitivity is increasing in baseline electricity consumption. As mentioned above, this is not surprising, as higher baseline consumption implies both a larger scope and incentives for consumption adjustment.
We also estimate the salience parameter θ to be decreasing in baseline electricity consumption, which is not surprising, either: The more biased a household is, the higher we would expect pre-treatment electricity consumption. Appendix Table 1.G.1 reports the p-values for our hypothesis tests of equality of the θ-estimates. Both the difference between the estimated θ for the high- and low-, as well as the difference in the estimated salience parameter for the high- and medium baseline

Table 1.5. GMM Heterogeneity Results

		Baseline Consumption			IHD-interaction		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Overall	High	Medium	Low	High	Medium	Low
θ	0.681***	0.439***	0.884***	1.179***	0.559***	0.645***	0.793***
	(0.072)	(0.141)	(0.071)	(0.103)	(0.133)	(0.122)	(0.135)
ψ_{11}^{FB}	-4.370***	-5.548***	-4.743***	-2.486***	-5.043***	-4.287***	-3.499***
	(0.442)	(1.038)	(0.645)	(0.505)	(0.836)	(0.720)	(0.610)
ψ_{21}^{FB}	-0.660***	-0.523*	-0.872***	-0.429***	-0.599***	-0.773***	-0.601***
	(0.107)	(0.290)	(0.122)	(0.074)	(0.159)	(0.229)	(0.139)
ψ_{22}^{FB}	-0.703	-1.477	2.068	-0.321	-1.105	-0.732	-2.244
	(0.659)	(1.019)	(2.259)	(0.713)	(0.903)	(1.060)	(1.858)
p-value "$\theta = 1$":	0.000	0.000	0.104	0.082	0.001	0.004	0.124
Hansen's J χ^2	2.718	3.038	0.481	2.809	1.725	0.342	3.719
p-value Hansen's J:	0.257	0.219	0.786	0.245	0.422	0.843	0.156
Observations	163'028	53'256	55'796	53'976	62'555	56'569	43'546
Clusters	907	302	304	301	303	312	289

Notes: This table presents the results for the salience parameter θ and the price sensitivities ψ from the GMM estimation based on the moment conditions that underlie the GMM framework, 1.28 and 1.29. The GMM estimation results in columns (2) - (7) are based on the tercile splits in the distributions of average baseline consumption and the number of interactions with the IHDs during the first 30 days after IHD-installation, as depicted in Figures 1.2 and 1.A.3. To estimate θ, a weighted average of different prices (depending on the electricity product) was used to approximate p_1 and p_2. A Wald-test tests whether the estimates for θ are significantly different from 1. Hansen's J statistic is used to test the validity of the overidentifying restrictions in a GMM model. All standard errors (in parentheses) are clustered at the household level. In Table 1.G.1 in the appendix, we test for the equality of θ across baseline consumption- and IHD-interaction cohorts.
*** $p < 0.01$ ** $p < 0.05$ * $p < 0.1$

households are significant at the 1 percent level, respectively.
Perhaps the most surprising result is our estimate for the salience parameter θ for low baseline households, as it is larger than one, indicating that low baseline households pay disproportionately more attention to their monetary electricity costs and exhibit a lower price sensitivity after IHD installation than before. This may appear surprising. Our theoretical framework is intentionally held very simple in order to make it as easy as possible to track the mechanisms that determine energy consumption and price responses.[19] The application process for the EWZ, however, provides starting points for more complex mechanisms. Like many other electricity providers, the EWZ offers an electricity costs calculator when applying for an electricity product, based on household size, tenancy status etc. This sets up a reference value for electricity costs and provides several potential starting points for mental accounting and reference dependence. At the end of the year, actual electricity consumption is compared to the reference value. Households thus set up a mental account for electricity costs and a reference point for consumption. Exceeding the reference point triggers loss aversion and provides incentives to conserve energy, whereas staying below it gives incentives to consume more energy to fully exploit the mental account for energy costs. As this possibility is not featured in our framework and beyond the scope of this paper, we will not further evaluate at this point.
Using equation (1.11) to recover the pre-treatment price sensitivities ψ^*_{11}, we calculate -2.474, -4.251 and -2.970 for high-, medium-, and low baseline households, respectively. We cannot explain why medium baseline households exhibit a substantially higher pre-treatment price sensitivity than high- and low baseline consumption households, but we can see that pre-treatment price sensitivity does not seem to be systematically related to baseline electricity consumption. However, after making households aware of their electricity consumption and the monetary savings potential of the Saturday/Sunday-pricing mechanism, high-baseline households (who show the highest savings potential) react the strongest, and low-baseline households react the weakest to the pricing mechanism. The heterogeneity in our results shows that some implications of our theoretical framework, such as the constant derivatives, are oversimplifications that are needed for identification, but do not perfectly reflect reality. For the sake of simplicity and identification, we assume a quadratic utility function and point out the potential shortcomings using the heterogeneity analysis.
Columns (5), (6) and (7) show the results from the IHD-usage tercile split. For this we split up the sample into households that interacted with the IHD by more or less than the 33 and the 66 percentile during the first 30 days after IHD-installation, according to Figure 1.A.3. By "IHD-interaction" we refer to the usage of an IHD, namely, by looking at current or past consumption, checking up on the (self-set)

19. Including reference-dependent utility (Kőszegi and Rabin (2006)) is generally possible, but would make the model too convoluted and is not the focus for this paper.

goal achievement progress or going into the "settings"-menu.
We first observe that our estimates for the salience parameter θ differ for high-, medium- and low-interaction households, though this difference is not statistically significant. We find that strongly biased households used the IHDs more than less biased households, indicating that biased households seek to "correct" their bias by using the IHD more frequently than their less biased counterparts and showing that households are - to some extend - aware of their own salience bias. IHD-interaction appears to be a driver of post-feedback price sensitivity. We estimate households that interacted with the IHDs a lot to be the most post-feedback price sensitive, whereas pre-feedback price sensitivity does not differ much across the three subgroups, with $\psi^*_{11} = -2.837$, -2.775 and -2.949, respectively for high-, medium- and low-interaction households. Our results indicate that biased households actively use the IHD to resolve the salience bias and that post-feedback price sensitivity is increasing in IHD-interaction. One might argue that households that actively sought out to use the IHD are also more determined to exploit the EWZ pricing structure, but note that we do not find large differences in the price sensitivities between households before the IHDs were installed.
So far, we have shown that IHD-interaction and pre-treatment baseline consumption are important mediators of the (cost) effectiveness of feedback-provision. However, it is likely that both characteristics are functions of even more fundamental sources of heterogeneity, such as household characteristics. In Appendix 1.D, we thus also analyze the underlying heterogeneity with respect to household income and the education level of the household head.

1.5 Consumer Surplus Implications

We now turn to the surplus implications of the feedback intervention, especially with respect to households' ability to adapt to the pricing schedule after being able to monitor their own consumption. Before diving into the analysis, an important remark is in order: The main motivation for this paper is not to analyze individual consumer surplus, but rather to quantify the salience bias and to show that feedback interventions are an effective tool to increase electricity price sensitivity and do not negatively affect the consumer surplus. Additionally, we want to show that the overall welfare gains through more flexible consumption adaption are not weighted down by individual surplus losses. As we will see in this section, households significantly save money by reducing their electricity consumption following feedback introduction, but experience much smaller gains in terms of consumer surplus. As suggested earlier this does not mean that feedback provision fails in terms of surplus gains. Individual surplus gains are a welcome side-effect of the more flexible energy adjustment of households, but not the main driver of the policy implications of this paper. As mentioned in Section 1.2.4, our theoretical framework allows us to calculate the surplus effects of IHD introduction based on our estimates. We can now plug in the estimated parameters from Table 1.3 in equation (1.14). A major advantage of the utility function we chose for our framework is the fact that it is linear in electricity costs. Hence, individual utility and consumer surplus are expressed in monetary units. Using the Delta-Method to recover the necessary statistics from our estimation results, we can calculate the implications of IHD-installation as follows:

$$W(1) - W(\theta) \;=\; -\Delta x_1^* \frac{(p_1 \Delta x_1^* + p_2 \Delta x_2^*)}{2\left(p_1 \frac{\partial x_1^{FB}}{\partial p_1} + p_2 \frac{\partial x_2^{FB}}{\partial p_1}\right)} \;\approx\; 0.011 \;\text{ CHF per day}$$

As consumption is defined on a daily level, this denotes the daily surplus gains of a household following IHD-installation. The total welfare gains accumulate to ≈ 4 CHF per year. Two remarks are in order: First, as announced earlier, observe that this number differs from pure cost savings, which are

$$\Delta x_1^* \cdot p_1 + \Delta x_2^* \cdot p_2 \;\approx\; 0.074 \;\text{ CHF per day}$$

Which amounts to 27 CHF per year. Though still positive, the surplus gains from providing households with feedback are substantially lower than the pure consumption savings, as households consume less energy and thus derive less utility from energy consumption. This is a substantial consumer surplus driver we must not ignore. However, remember that the main motivation for our analysis was not to evaluate individual consumer surplus as defined in Section 1.2.4. Instead, we try to adapt total energy consumption to the energy supply to relieve electricity plants and grids at peak consumption hours, thus increasing overall welfare. Hence, the main objective of our welfare analysis is to show that households' consumer surplus is not negatively affected by the intervention despite lower energy consumption.

1.5.1 Heterogeneity in Consumer Surplus Effects

How differently is a household's surplus affected by the intervention, conditional on household characteristics? If we were able to quantify the surplus effect on different subsets of households, we could be able to tailor the pricing schedule to the households in order to maximize both the pricing effect and surplus gains.[20] Table 1.6 presents the monetary savings and consumer surplus gains for each subset of households analyzed in Section 1.4.7 based on equation (1.14).

20. Note that in this section, unless otherwise specified, the "treatment effect of IHD-introduction" refers to the treatment effect on DT-consumption, as this is the main driver of our welfare analysis.

Table 1.6. Daily Consumer Surplus Gains

	(1)	(2)
	Monetary Savings	CS Gains
Overall	0.074***	0.011**
	(0.017)	(0.006)
High Baseline Consumption	0.160***	0.067*
	(0.037)	(0.036)
Medium Baseline Consumption	0.026	0.002
	(0.018)	(0.002)
Low Baseline Consumption	-0.024*	0.001
	(0.013)	(0.001)
High IHD usage	0.116***	0.032*
	(0.030)	(0.018)
Medium IHD usage	0.083	0.014
	(0.031)	(0.012)
Low IHD usage	0.043	0.003
	(0.026)	(0.004)

Notes: The table presents approximations of daily monetary savings and daily gains in consumer surplus using the Delta Method based on equation (1.14) and the estimates from Tables 1.3 and 1.5 (in CHFs per day).
*** $p < 0.01$ ** $p < 0.05$ * $p < 0.1$

As expected, we observe the consumption savings to be increasing in the treatment effect of IHD-introduction and the estimated salience bias: The high baseline consumption- and high IHD-interaction households drastically adapt their consumption behavior after being introduced to consumption feedback and hence, experience the largest savings. If the treatment effect is negative (i.e. if feedback increases consumption), the consumption savings are negative, as can be seen for low baseline consumption households. The consumer surplus gains from IHD-introduction are also increasing in the treatment effect and post-feedback price sensitivity, but exhibit a much smaller magnitude, with the surplus gains being significantly positive only for high IHD interaction and high baseline consumption households. As pointed out in Section 1.2.4, we observe that the weighted sum of the surplus gains exceeds the surplus gains for the whole sample:

$$0.33 \cdot 0.067 + 0.33 \cdot 0.002 + 0.33 \cdot 0.001 \approx 0.023 > 0.011$$

1.6 Discussion and Conclusion

The externalities of excessive fossil and nuclear energy production have been demonstrated in the form of pollution, climate change and nuclear disasters and clearly show the necessity for a more structured approach to sustainable energy management. In order to facilitate the transition towards renewable energy sources, it will be necessary to reduce aggregate energy consumption and to flexibly adapt energy demand to the contemporaneous energy supply. By extending the strategy of a randomized field experiment conducted in Zurich, Switzerland, and by exploiting the nonlinear pricing mechanism used by the local electricity provider, we are able to quantify how much attention households pay towards energy prices. Additionally, we are able to quantify the conservation potential of electricity price variation in combination with the installation of smart meters and In-Home-Displays.

Our results indicate that, on average, consumers perceive less than 70 percent of the electricity costs they actually incur. In line with current literature, we argue that this salience bias arises from the discrepancy in the salience of energy consumption and energy costs, which are usually incurred with delay. We find that, on average, IHD-introduction decreases daytime consumption on Saturdays by 0.32 kWh (≈ 6 percent of average energy consumption during the daytime tariff), on Sundays by 0.21 kWh (≈ 4 percent) and on weekdays by 0.24 kWh (≈ 5 percent) and has a somewhat ambiguous, though insignificant effect on nighttime electricity consumption. This indicates that feedback induces households to exploit the existing pricing schedules. Additionally, we find that IHD-introduction increases price sensitivity: Before intervention, a price increase of 1 CHF decreased daytime energy consumption by 3.1 kWhs per day on average, and after intervention by 4.4 kWhs per day.

Our findings are subject to substantial heterogeneity in household characteristics, with our heterogeneity analysis suggesting the salience bias to be stronger for households with high baseline consumption, low education level and low income. Our finding that IHD interaction indeed mediates treatment and price effects confirms earlier results that consumers need to engage with the IHDs to realize the best outcomes. Conversely, the results are not rendered insignificant if they do not.

The results hold important implications for both policymakers as well as future research: When introducing taxes or manipulating the energy price to reduce energy consumption and evaluating welfare effects, policymakers have to acknowledge the importance of salience of electricity prices. As Chetty, Looney, and Kroft (2009), Finkelstein (2009) and Taubinsky and Rees-Jones (2017) point out: for most cases of taxation, consumer inattention and the resulting low price elasticity is actually desirable: Basic economics tells us that low price elasticities minimize the deadweight loss inherent to taxation. But corrective taxes, such as carbon taxes on electricity consumption, are not primarily aimed at generating tax revenue, but decreasing demand and hence, consumption (Allcott, Lockwood, and Taubinsky (2019)). Keeping the balance between minimizing welfare losses and regulating energy consumption demands a deep understanding of consumer inattention and price elasticities. This study provides a starting point on this front by quantifying the level of inattention towards electricity prices. Even further, we are able to quantify consumer surplus effects by pointing out the discrepancy to simple conservation effects. Our straightforward framework allows future research to easily replicate and extend our framework using different datasets and to compare the results to ours. Differences in daytime and nighttime prices are common in most countries - it is likely that other electricity providers use a similar pricing mechanism as the EWZ. Hopefully, the simplicity of our design initiates further scientific work that investigates the importance of feedback on the electricity demand sensitivity with respect to the price. The most compelling extension to our theoretical framework and the empirical setting would be a three-good framework, in which electricity consumption over the day is further split up into different hours of the day. Using this approach, it would be possible to further shed light on our surprising findings regarding the sign of the derivative of daytime electricity consumption with respect to the nighttime electricity price. The nature of the price variation in this setting does not allow for such an approach.

The data structure and the simplicity of our structural approach have, of course, some shortcomings: As we see in Figure 1.4, Saturday and Sunday load profiles differ from regular weekday load profiles. By restricting our analysis to Saturdays and Sundays we ignore these differences. This may raise questions regarding the external validity of our design, as we cannot easily extend our results to weekdays. However, by showing that the treatment effect of IHD-introduction on weekdays in extremely similar to the treatment effect on Saturdays we believe that our results can - at least to some extend - be generalized to weekdays. Our robustness check confirms this: The price variation between Saturdays and Sundays does not allow for a GMM approach (as Sundays and Mondays are too different to be compared directly), but a Difference-in-Differences framework can be imposed on the data and the Sunday-Monday price variation can be used to estimate the salience parameter θ. As Table 1.F.3 reveals, this does not change our estimate for the salience parameter substantially. Our description of the salience bias in this setting remotely resembles the *Present Bias* (Laibson (1997)): Consumers disproportionately undervaluing future costs could also be explained by time-biased preferences and hyperbolic discounting. While our framework does not distinguish between the two possibilities, we believe the feedback intervention to only have a sizable effect on electricity consumption if not time-biased preferences, but limited salience to future costs is the main driver of excessive pre-treatment energy consumption. This paper contributes to the literature on interventions used in environmental psychology (Steg and Vlek (2009)) and aimed at pro-environmental behavior. Future research could further pursue the welfare implications of our empirical strategy and our theoretical framework. The possible extensions to our framework and strategy include taxation effects under salience bias and the development of treatment effects and load shifting over time. As electricity pricing mechanisms are nonlinear in many countries, our analysis can also be conducted on a variety of already existing datasets that - much like in the presented case for Degen et al. (2013) - are generated to examine the consumption outcomes after feedback provision, without taking a closer look at price effects.

Appendix 1.A Additional Figures

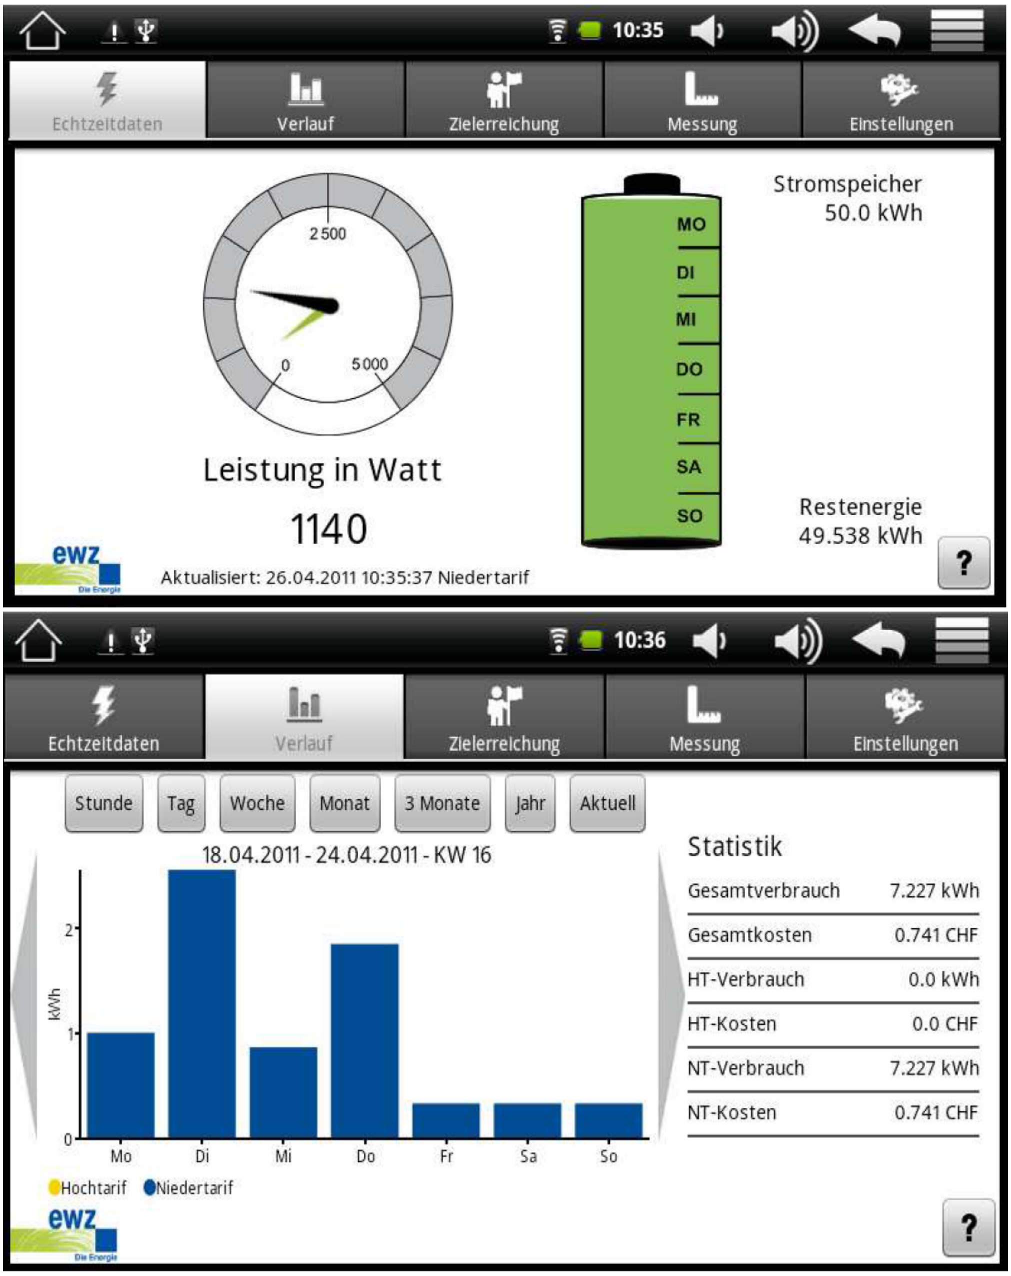

Figure 1.A.1. In-Home-Display Interface

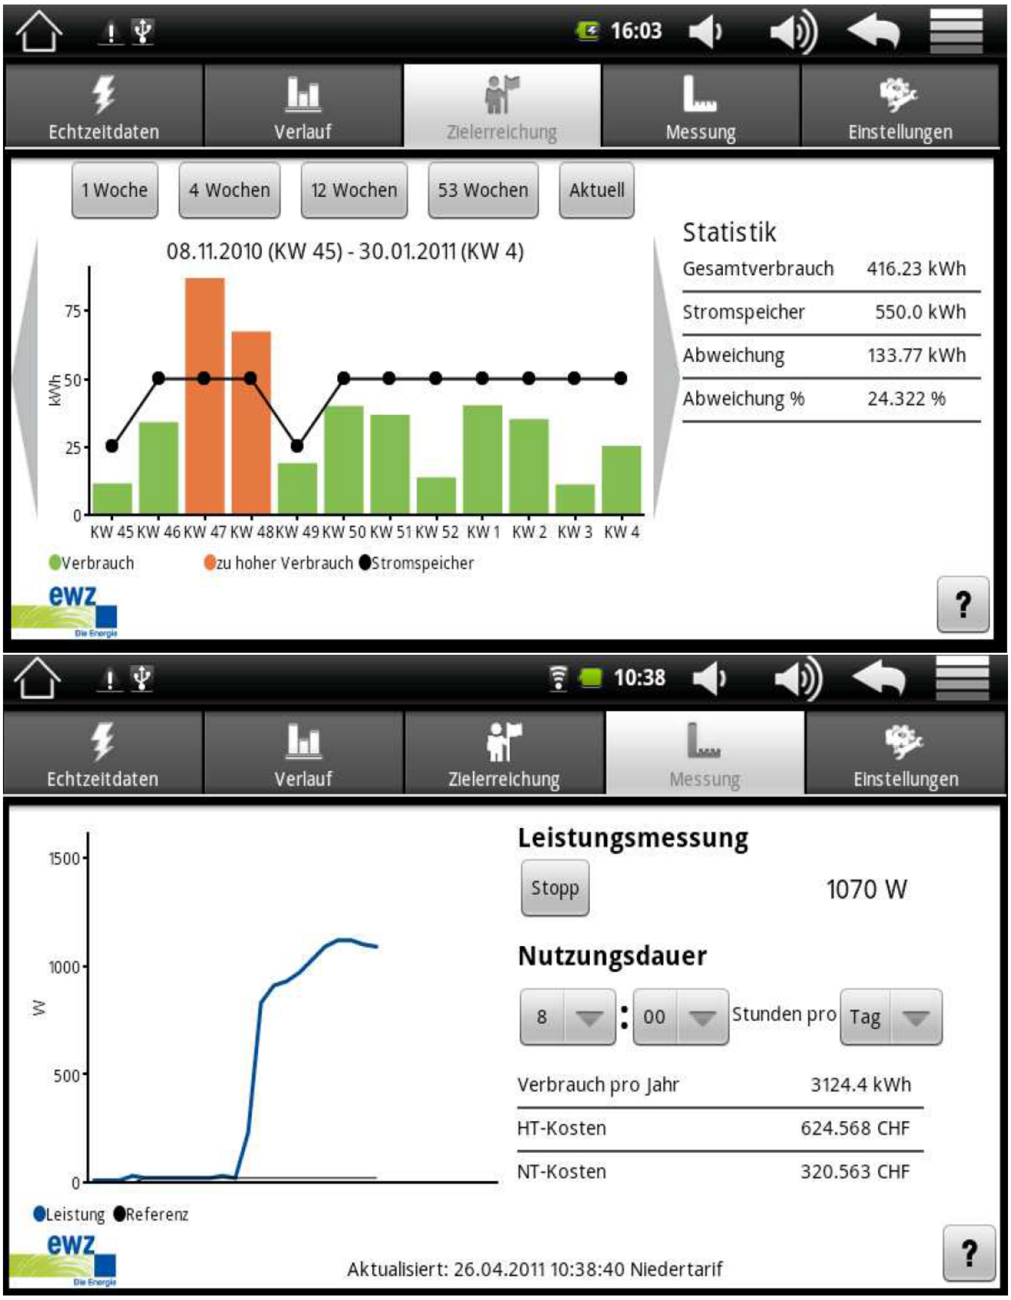

Figure 1.A.2. In-Home-Display Interface

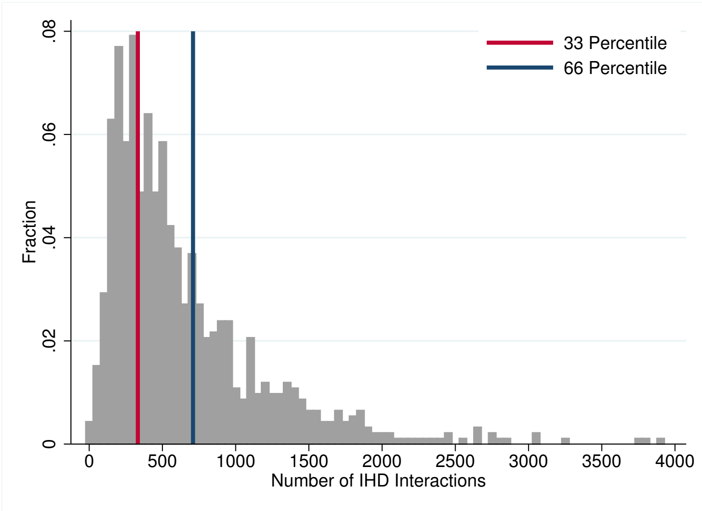

Figure 1.A.3. Distribution of the Number of IHD Interactions

This figure presents the distribution of the number of IHD-interactions per household during the first 30 days after IHD-installation. For clarity reasons, we exclude households with more than 4000 interactions.

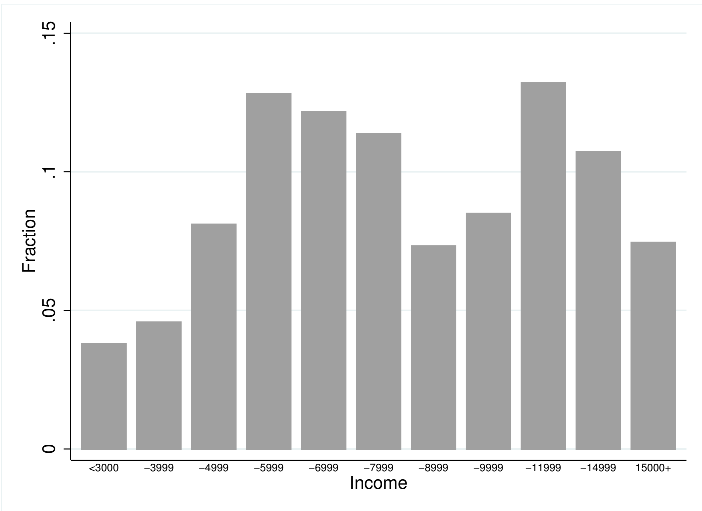

Figure 1.A.4. Distribution of Household Income

This figure presents the distribution of household income. Note that -3999 is short for 3000-3999, -4999 is short for 4000-4999 and so on. For the correlation Table 1.4, we assign to each income category a value between 1 and 11.

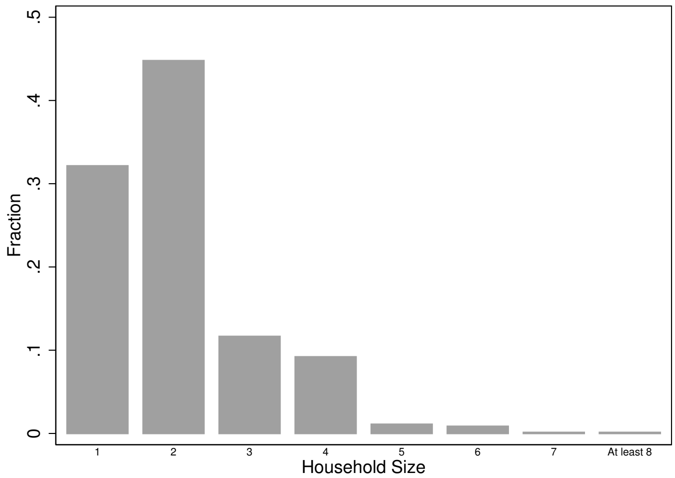

Figure 1.A.5. Distribution of Household Size

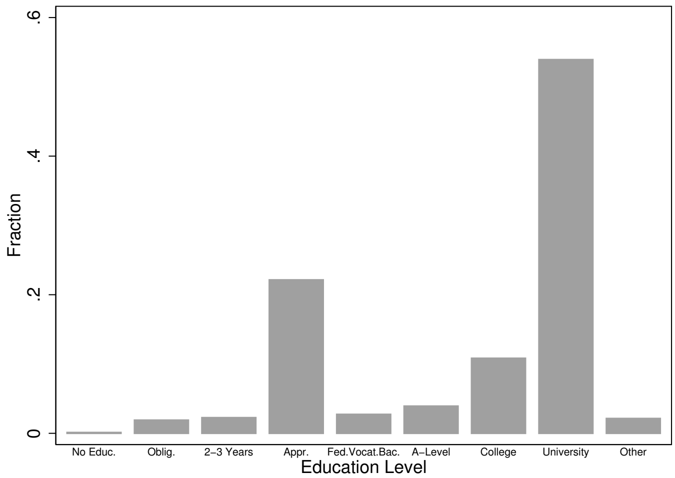

Figure 1.A.6. Distribution of Education Level

This figure presents the distribution of the highest education level of the household heads. The included education levels are: *No Education, Obligatory Schooling, 2-3 Years General Schooling, Apprenticeship, Federal Vocational Baccalaureat, A-Level, College, University and Others*. For the correlation Table 1.4, we assign to each education category (except *other*) a value between 1 and 8 according to the education hierarchy displayed in this graph.

Price in CHF per kWh — Sunday

Product	NT	DT
ökopower	0.1459	0.1459
naturpower	0.1026	0.1026
mixpower	0.0972	0.0972
solartop	0.7022	0.7022

Price in CHF per kWh — Other Days

Product	NT	DT
ökopower	0.1459	0,243
naturpower	0.1026	0,1999
mixpower	0.0972	0,1945
solartop	0.7022	0,7022

Figure 1.A.7. Electricity Products and Prices

Electricity products and their respective prices on Sundays and every other day of the week. Note that we excluded households with the *SolarTop*-tariff from our analyses as their pricing schedule is different from the standard pricing schedule. Additionally, they make up only a small fraction of our total sample (0.35%).

Appendix 1.B Robustness Check: Analysis Without Cutoff Hour

Table 1.B.1. Average Treatment Effect for Each Day of the Week and Time of the Day Without the DT Cutoff Hour (9:30 - 10:30 p.m.)

Dependent Variable	(1) Consumption	(2) Consumption	(3) Consumption	(4) Consumption	(5) Consumption	(6) Consumption
Treatment Effect Saturday DT	-0.298***					
	(0.071)					
Treatment Effect Saturday NT		-0.020				
		(0.022)				
Treatment Effect Sunday DT			-0.193***			
			(0.066)			
Treatment Effect Sunday NT				-0.032		
				(0.021)		
Treatment Effect Weekday DT					-0.230***	
					(0.047)	
Treatment Effect Weekday NT						0.016
						(0.017)
Individual-FEs	Yes	Yes	Yes	Yes	Yes	Yes
Week-FEs	Yes	Yes	Yes	Yes	Yes	Yes
Observations	41'577	41'617	41'446	41'506	208'033	208'269
R^2_{within}	0.055	0.008	0.057	0.009	0.040	0.006
$R^2_{between}$	0.007	0.000	0.009	0.000	0.007	0.001
$R^2_{overall}$	0.028	0.003	0.029	0.003	0.021	0.002
Clusters	941	941	940	940	942	942

Notes: This table presents results from the DiD-regressions based on equations (1.16)-(1.21). Dependent variable: Daily DT or NT electricity consumption, respectively. In all estimations we exclude observartions during the cutoff hour 9:30 - 10:30 p.m. The estimated coefficients denote mean deviations from pre-treatment baseline consumption, based on indicators for each combination of day-of-the-week and DT/NT-tariff. The control group in this framework consists of households with an installed smart meter, but no feedback through IHDs. Each column reports the mean deviation (= average treatment effect) for each day-of-the-week based on equations (1.16) - (1.21). Each treatment effect was estimated in a separate DiD regression. The control group in each specification consists of households with an installed smart meter, but which do not receive feedback through an IHD, yet. All specification include week-of-sample- and household fixed effects. Standard errors (in parentheses) are clustered at the household level.

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

Table 1.B.2. GMM Results without the DT Cutoff Hour (9:30 - 10:30 p.m.)

	Estimate
θ	0.680***
	(0.074)
ψ_{11}^{FB}	-4.149***
	(0.427)
ψ_{21}^{FB}	-0.425***
	(0.079)
ψ_{22}^{FB}	-0.553
	(0.564)
p-value $\theta = 1$:	0.000
Hansen's J χ^2	2.633
p-value Hansen's J:	0.268
Observations	164'331
Clusters	918

Notes: This table presents the results for the salience parameter θ and the price sensitivity ψ from the GMM estimation based on the moment conditions that underlie the GMM framework, 1.28 and 1.29, excluding observation during the DT cutoff hour 9:30 - 10:30 p.m. All standard errors (in parentheses) are clustered at the household level. To estimate θ, a weighted average of different prices (depending on the electricity product) was used to approximate p_1 and p_2. A Wald-test tests whether the estimates for θ are significantly different from 1. Hansen's J statistic is used to test the validity of the overidentifying restrictions in a GMM model.

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

Appendix 1.C Average Treatment Effect over the Day

In this section, we present reduced form evidence on the treatment effect of IHD-installation for each hour of the day. That is, we estimate the Difference-in-Differences framework presented in equation (1.C.1) separately for Saturdays, Sundays and weekdays. The control group consists of households with an installed smart meter, but who did not receive feedback through an IHD, yet. The outcome variable is hourly electricity consumption. We include household and week-of-sample fixed effects as well as hour-of-the-day fixed effects. The δ-coefficients then deliver the average treatment effects for each hour of the day. The results closely resemble our results in Figure 1.6, in that, following IHD-installation, we observe the largest reductions in consumption during the peak hours in the evening and on Saturdays (not on Sundays, when consumption is highest). Additionally, we find reductions of similar magnitude on weekdays.

$$y_{ith} = \alpha_i + \beta_{w(t)} + \sum_{x=0}^{23} \delta_x^{after} \mathbf{I}_{(h=x \,\&\, after)} + \sum_{x=0}^{23} \gamma_x \mathbf{I}_{(h=x)} + \epsilon_{ith} \tag{1.C.1}$$

The estimated coefficients are displayed in Figure 1.C.1c.

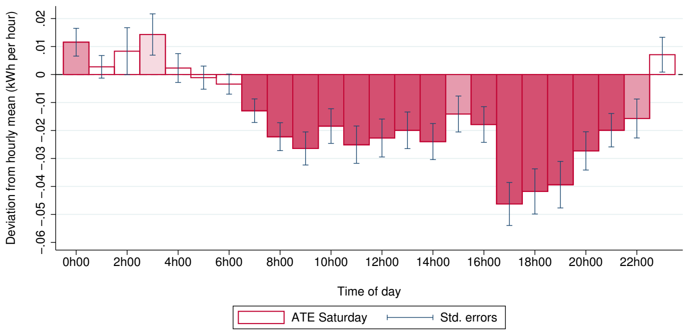

(a) Average Treatment Effect over Saturdays

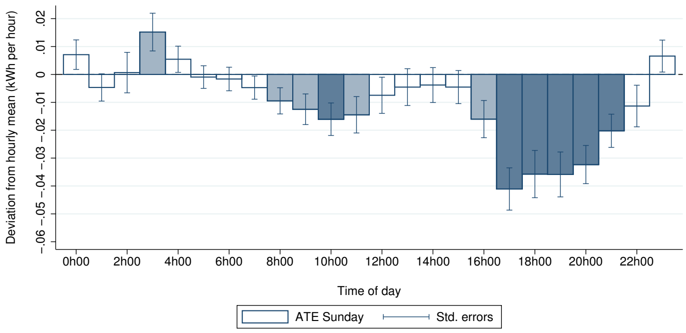

(b) Average Treatment Effect over Sundays

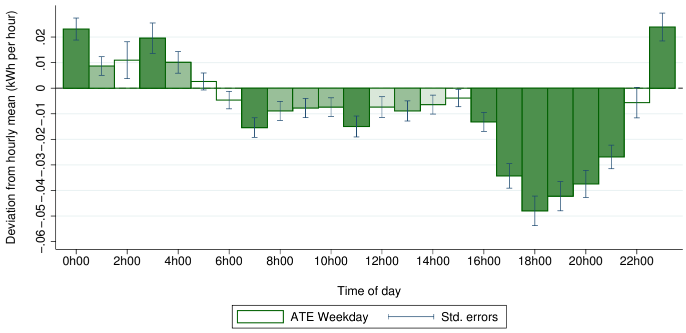

(c) Average Treatment Effect over weekdays

Figure 1.C.1. Averagte Treatment Effect over the Day on Different Days of the Week

This figure presents the estimated treatment effect for each hour of the day for Saturdays, Sundays, and weekdays respectively, based on equation (1.C.1). The outcome variable is hourly electricity consumption. We include household and week-of-sample fixed effects as well as hour-of-the-day fixed effects. The control group consists of households with smart meter, but no feedback through IHDs, yet. Shaded bars indicate statistical significance at the 10-, 5-, and 1-percent level, respectively. Whiskers denote the standard error (clustered at the household level). This graph was inspired by Andor, Gerster, and Goette (2020)

Appendix 1.D Heterogeneity with Respect to Income and Education Level

Table 1.D.1 reports our results from the household income heterogeneity analysis. Figure 1.A.4 presents the income distribution across all households in our sample. Looking at Table 1.4, note that income is strongly correlated with initial baseline consumption, household size and education level. Column (1) in Table 1.D.1 reports the GMM results for the overall sample; columns (2), (3) and (4) report the results for low-, medium- and high income households. We estimate low-income households to be the most biased households in our sample. The difference in the estimate for the salience parameter between low- and medium income households is mostly driven by the large difference in pre-treatment price sensitivity. After being introduced to the IHDs, we observe little differences in the price sensitivities of both subsets, indicating that low-income households show little sign of adaption to the pricing schedule pre-treatment, but adapt their consumption behavior once they are made aware of the savings potential. Interestingly, the largest pre- and post feedback price sensitivities can be found in the high income households. Taking a look at the correlation structure between baseline consumption and income helps us to better understand the results in Table 1.D.1.

Table 1.D.2 presents our estimation results based on the education level of the survey respondent. Figure 1.A.6 displays the distribution of the education levels of the survey respondents. Table 1.4 reveals that education level is negatively correlated with baseline consumption. Note that *education* is a categorical variable, and including it in a correlation table implicitly requires that we can form a hierarchy consisting of the different education levels. We assume this hierarchy to correspond to the order in which the education levels are reported in Figure 1.A.6 (excluding "Others"). Column (2) presents the results for households in which the survey respondent stated that they finished an apprenticeship, the Federal Vocational Baccalaureat, A-Level or College; Column (3) presents the results for university-educated households. In the following, we will refer to them as *low* and *high* education households. We do not observe any significant differences in the post-feedback price sensitivities between low and high education households. However, we find that low-education households are stronger biased than high education households, with the difference being significant at the 10 percent level ($p = 0.05$)

Table 1.D.1. GMM Heterogeneity Estimation Results Based on Income

	(1)	(2)	(3)	(4)
		Income Cohort		
	Overall	< 6000CHF/Month	$6000 - 8999$CHF/Month	≥ 9000CHF/Month
θ	0.681***	0.409**	0.848***	0.836***
	(0.072)	(0.204)	(0.240)	(0.079)
ψ_{11}^{FB}	-4.370***	-2.801***	-2.401***	-7.503***
	(0.442)	(0.552)	(0.926)	(0.918)
ψ_{21}^{FB}	-0.660***	-0.402***	-0.286	-1.244***
	(0.107)	(0.136)	(0.204)	(0.224)
ψ_{22}^{FB}	-0.703	0.091	0.163	-0.665
	(0.659)	(0.482)	(2.336)	(2.705)
p-value "$\theta = 1$":	0.000	0.004	0.528	0.037
Hansen's J χ^2	2.718	1.728	1.656	0.756
p-value Hansen's J:	0.257	0.422	0.437	0.685
Observations	163'028	40'242	43'550	54'179
Clusters	907	219	233	301

Notes: This table presents the results for the salience parameter θ and the price sensitivities ψ from the GMM estimation based on the moment conditions that underlie the GMM framework, 1.28 and 1.29. The GMM estimation results in columns (2) - (4) are based on household income cohorts. In column (2), we present the results for households earning less than 6000 CHF per month, in column (3), only households earning between 6009 - 8999 are included, and in column (4), we only include households earning 9000 CHF per month or more. To estimate θ, a weighted average of different prices (depending on the electricity product) was used to approximate p_1 and p_2. A Wald-test tests whether the estimate for θ is significantly different from 1. Hansen's J statistic is used to test the validity of the overidentifying restrictions in a GMM model. In Table 1.G.1 we test for the equality of θ across income cohorts. All standard errors (in parentheses) are clustered at the household level *** $p < 0.01$ ** $p < 0.05$ * $p < 0.1$

Table 1.D.2. GMM Heterogeneity Results Based on Education

	(1)	(2)	(3)
		Education Cohort	
	Overall	"Low" Education	"High" Education
θ	0.681***	0.526***	0.844***
	(0.072)	(0.137)	(0.088)
ψ_{11}^{FB}	-4.370***	-4.020***	-4.488***
	(0.442)	(0.742)	(0.603)
ψ_{21}^{FB}	-0.660***	-0.338**	-0.969***
	(0.107)	(0.152)	(0.162)
ψ_{22}^{FB}	-0.703	-0.391	0.115
	(0.659)	(0.523)	(2.113)
p-value "$\theta = 1$":	0.000	0.000	0.075
Hansen's J χ^2	2.718	3.323	2.511
p-value Hansen's J:	0.257	0.190	0.285
Observations	163'028	67'909	80'481
Clusters	907	365	445

Notes: This table presents the results for the salience parameter θ and the price sensitivities ψ from the GMM estimation based on the moment conditions that underlie the GMM framework, 1.28 and 1.29. Column (1) presents the results for households in which the survey respondent stated that they finished an apprenticeship, the Federal Vocational Baccalaureat, A-Level or College, Column (2) presents the results for university-educated households. Households of the *other*-category are excluded from the analysis. To estimate θ, a weighted average of different prices (depending on the electricity product) was used to approximate p_1 and p_2. A Wald-test tests whether the estimate for θ is significantly different from 1. Hansen's J statistic is used to test the validity of the overidentifying restrictions in a GMM model. In Table 1.G.1 we test for the equality of θ across education cohorts. All standard errors (in parentheses) are clustered at the household level.

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

Appendix 1.E Average Treatment Effect: Heterogeneity

Table 1.E.1 presents the average treatment effects of IHD-introduction on Saturday DT electricity consumption for the three subsamples of high baseline, medium baseline and low baseline households. We applied the same estimation equation (1.16) we used to receive the estimates in tables 1.2 and 1.B.1. We observe significant reductions following IHD-introduction only for the high baseline households. We can also see that low baseline households increased electricity consumption. Although this effect is not significantly different from zero, we see that it is significantly different from the average treatment effect on the high baseline households. We also test for equality among the treatment effects for each baseline consumption tercile and report the corresponding p-values. Finally, we run the same analysis for a tercile split on the number of IHD-interactions during the first 30 days after IHD-installation and find that households in all terciles significantly reduced consumption, but even more so in the high-IHD-interaction tercile (though the difference to the remaining terciles is not statistically significant).

Table 1.E.1. Average Treatment Effect for Each Baseline Tercile

	Estimate
Treatment Effect High Baseline	-0.824***
	(0.136)
Treatment Effect Medium Baseline	-0.119
	(0.087)
Treatment Effect Low Baseline	0.045
	(0.088)
***p*-values**	
High vs. Medium	0.000
Medium vs. Low	0.109
High vs. Low	0.000
Individual-FEs	Yes
Week-FEs	Yes
Observations	41’180
R^2_{within}	0.057
$R^2_{between}$	0.222
$R^2_{overall}$	0.005
Clusters	920

Notes: This table presents the average treatment effects of IHD-introduction on Saturday DT electricity consumption for the three subsamples of high baseline, medium baseline and low baseline households. Each treatment effect was estimated in one DiD regression based on equation (1.16). The tercile split is displayed in figure 1.2. Standard errors (in parentheses) are clustered at the household level. The p-values stem from simple Wald-tests testing the equality of the estimated average treatment effects of IHD-installation.
*** $p < 0.01$ ** $p < 0.05$ * $p < 0.1$

Table 1.E.2. Average Treatment Effect for Each IHD-interaction Tercile

	(1) Estimate
Treatment Effect High IHD Usage	-0.426***
	(0.125)
Treatment Effect Medium IHD Usage	-0.285***
	(0.107)
Treatment Effect Low IHD Usage	-0.276***
	(0.091)
***p*-values**	
High vs. Medium	0.342
Medium vs. Low	0.940
High vs. Low	0.285
Individual-FEs	Yes
Week-FEs	Yes
Observations	41’180
R^2_{within}	0.054
$R^2_{between}$	0.002
$R^2_{overall}$	0.024
Clusters	920

Notes: This table presents the average treatment effects of IHD-introduction on Saturday DT electricity consumption for the three subsamples of high- medium, and low IHD-interaction. Each treatment effect was estimated in one DiD regression based on equation (1.16) and a tercile split on the number of IHD interactions during the first 30 days after IHD installation. The tercile split is presented in figure 1.A.3. Standard errors (in parentheses) are clustered at the household level. The *p*-values stem from Wald-tests testing the equality of the estimated average treatment effects of IHD-installation.
*** $p < 0.01$ ** $p < 0.05$ * $p < 0.1$

Appendix 1.F Robustness Check: Difference-in-Differences Setup

Using a difference-in-differences setup, we can check whether our results are robust to different specifications. Note that the difference-in-differences analysis was conducted before the GMM analysis was implemented and provides the starting values for the parameters identified in Section 1.4.6.

Using the restriction of the data to Saturdays and Sundays, the random assignment to the control group and the assumptions of constant derivatives and full resolution of the salience bias under feedback from the theoretical framework, we can estimate the necessary statistics in the fixed-effects model presented in equation (1.F.1).

$$\begin{aligned} y_{it\tau} &= \alpha_i + \omega_{w(t)} \\ &\quad + \sum_{j(t)\in\{Sa,Su\}} \sum_{\tau\in\{DT,NT\}} \sum_{\gamma\in\{before,after\}} \beta^{j\tau\gamma} D_{ij(t)\tau\gamma} - \beta_{Sa,DT,before} + \epsilon_{it\tau\gamma} \end{aligned} \tag{1.F.1}$$

That is, household i's consumption at date t at time $\tau \in \{DT, NT\}$ is determined by the day of the week $j(t)$, its respective pricing structure for the DT and NT (denoted by τ) and by the provision of feedback through IHDs (denoted by *before* or *after*). The idea behind equation (1.F.1) is as follows: $D_{ij(t)\tau\gamma}$ is an indicator, which is equal to 1 if an observation for household i falls on a certain day of the week $j(t)$ (*Sa* or *Su*), in a certain time tariff $\tau \in \{DT, NT\}$ and in the time *before* or *after* feedback, denoted by γ. Hence, for each household, there are $2^3 = 8$ possible combinations of weekday-time-feedback. For each combination, we include a dummy variable. To avoid perfect collinearity (because exactly one of these indicators will, by construction, be 1), we exclude the dummy for Saturday during the DT before ($Sa, DT, before$) intervention, as this will serve as the reference category in our regression. Finally, we include household fixed-effects α_i and week-of-the-year fixed effects $\omega_{w(t)}$ and cluster all standard errors on the household-level. Table 1.F.1 presents the results from the regression described above.

Before interpreting the results, we will translate the coefficients on the dummy-variables from equation (1.F.1) to fit our model. Using the fact that consumption on Saturday during the DT before intervention serves as the reference category, we see that the coefficient on $D_{i,Sa,DT,after}$ gives us the average treatment effect of providing feedback on Saturdays during the DT. In other words: the coefficient on $D_{i,SA,DT,after}$ gives us Δx_1^*. To see this, remember that the coefficients in Table 1.F.1 report mean deviations from the reference category, meaning that the coefficient on $D_{i,Sa,DT,after}$ denotes how much more or less energy the average household consumes on Saturday during the DT after intervention compared to the reference category *Saturday DT before* intervention.

Using the same logic, we can find the DT-price sensitivity of the average household

under feedback (ψ_{11}^{FB}). From the pricing mechanism, we already know $\Delta p_1 = 0.0973$ CHF. Δx_1^{FB} can be calculated as $\hat{\beta}^{Su,DT,after} - \hat{\beta}^{Sa,DT,after}$, meaning that the change from Saturday during the DT after the intervention to Sunday during the DT after intervention can be calculated as the difference between the coefficients on these dummies. Analogously, we can calculate Δx_2^{FB} as $\hat{\beta}^{Su,NT,after} - \hat{\beta}^{Sa,NT,after}$.
Note that in order to estimate θ from the data using equation (1.11), we have to estimate a nonlinear function of our estimated coefficients. We will thus make use of the Delta-Method.

Table 1.F.1. Main Results from Difference-in-Differences Specification

	Estimate
$\beta^{Sa,DT,after}$	-0.401***
	(0.066)
$\beta^{Sa,NT,before}$	-3.358***
	(0.093)
$\beta^{Sa,NT,after}$	-3.328***
	(0.096)
$\beta^{Su,DT,before}$	0.292***
	(0.056)
$\beta^{Su,DT,after}$	0.017
	(0.075)
$\beta^{Su,NT,before}$	-3.295***
	(0.095)
$\beta^{Su,NT,after}$	-3.264***
	(0.098)
Individual-FEs	Yes
Week-FEs	Yes
Observations	164'554
R^2_{within}	0.344
$R^2_{between}$	0.004
$R^2_{overall}$	0.234
Clusters	920

Notes: This table presents estimation results based on equation (1.F.1). The reported coefficients denote mean deviations from the reference category *Saturday DT before Intervention*. Standard errors (in parentheses) are clustered at the household level.
*** $p < 0.01$ ** $p < 0.05$ * $p < 0.1$

We can now calculate

$$\Delta x_1^{FB,p_1} = \hat{\beta}^{Su,DT,after} - \hat{\beta}^{Sa,DT,after} = 0.017 - (-0.401) = 0.418$$
$$\Delta x_2^{FB,p_1} = \hat{\beta}^{Su,NT,after} - \hat{\beta}^{Sa,NT,after} = -3.264 - (-3.328) = 0.064$$

Next, we can calculate the price derivatives of DT- and NT-consumption under feedback with respect to the DT-price:

$$\frac{\partial x_1^{FB}}{\partial p_1} = \frac{\Delta x_1^{FB,p_1}}{\Delta p_1} = \frac{0.419}{-0.0973} = -4.306 \tag{1.F.2}$$

$$\text{and} \quad \frac{\partial x_2^{FB}}{\partial p_1} = \frac{\Delta x_2^{FB,p_1}}{\Delta p_1} = \frac{0.064}{-0.0973} = -0.658 \tag{1.F.3}$$

Table 1.F.2 reports all approximations using the Delta Method. The salience parameter θ is reported twice, θ_1 reports the calculated salience parameter based on equation (1.8) while θ_2 is based on equation (1.11).

Table 1.F.2. Calculated Statistics from DiD-Setup

	Coefficient
Δx_1^*	-0.401***
	(0.066)
$\Delta x_1^{FB,p_1}$	0.419***
	(0.044)
$\Delta x_2^{FB,p_1}$	0.064***
	(0.011)
$\Delta x_1^{*,p_1}$	0.292***
	(0.056)
ψ_{11}^{FB}	-4.306***
	(0.456)
ψ_{21}^{FB}	-0.654***
	(0.111)
ψ_{11}^*	-3.001***
	(0.578)
θ_1	0.588***
	(0.076)
p-value $\theta_1 = 1$	0.000
θ_2	0.699***
	(0.117)
p-value $\theta_2 = 1$	0.010
N	164'554

Notes: This table presents the calculated statistics from the baseline difference-in-differences regression presented in Table 1.F.1 needed to calculate the salience parameter θ based on equations (1.8) and (1.11). To approximate these statistics, we used the Delta Method and the framework laid out earlier in this section.

Table 1.F.3. Results from Sunday-Monday Difference-in-Differences Specification

	Estimate
$\beta^{Su,DT,after}$	-0.285***
	(0.063)
$\beta^{Su,NT,before}$	-3.589***
	(0.092)
$\beta^{Su,NT,after}$	-3.560***
	(0.094)
$\beta^{Mo,DT,before}$	-0.574***
	(0.068)
$\beta^{Mo,DT,after}$	-0.959***
	(0.081)
$\beta^{Mo,NT,before}$	-3.715***
	(0.095)
$\beta^{Mo,NT,after}$	-3.672***
	(0.095)
θ	0.852***
	(0.080)
p-value $\theta = 1$	0.065
Week-FEs	Yes
Observations	164'165
R^2_{within}	0.342
$R^2_{between}$	0.003
$R^2_{overall}$	0.237
Clusters	919

Notes: This table presents estimation results based on equation (1.F.1). Reported coefficients denote mean deviations from the reference category *Sunday DT before Intervention*. Coefficients are estimated using OLS. The difference to the estimation above is that in this estimation, we only include observations on Sundays and Mondays and thus exploit the price difference during the DT between Sunday and Monday. As we cannot plausibly assume that Sunday and Monday are structurally similar enough to allow for for a direct comparison in energy consumption, we cannot use the identification of θ according to equation (1.8) for the calculation. Instead, θ is approximated using the Delta Method, based on equation 1.11. θ was approximated using the Delta Method. Standard errors (in parentheses) are clustered at the household level.
*** $p < 0.01$ ** $p < 0.05$ * $p < 0.1$

Appendix 1.G Hypothesis Tests

Table 1.G.1. *p*-values for Heterogeneity Analyses

Hypothesis Test *p*-values	
	p-value
	Baseline Consumption
$\theta_{High} = \theta_{Medium}$	0.005
$\theta_{Medium} = \theta_{Low}$	0.018
$\theta_{High} = \theta_{Low}$	0.000
	IHD-Interaction
$\theta_{High} = \theta_{Medium}$	0.635
$\theta_{Medium} = \theta_{Low}$	0.416
$\theta_{High} = \theta_{Low}$	0.216
	Income Cohort
$\theta_{High} = \theta_{Medium}$	0.165
$\theta_{Medium} = \theta_{Low}$	0.961
$\theta_{High} = \theta_{Low}$	0.052
	Education Level
$\theta_{High} = \theta_{Low}$	0.051

Notes: In this table, we display *p*-values for Wald tests of equality of coefficients according to the estimation results from Tables 1.5, 1.D.1 and 1.D.2. We test whether the estimate for the salience parameter θ differs between certain subgroups in our sample.

References

Abrahamse, Wokje, Linda Steg, Charles Vlek, and Talib Rothengatter. 2007. "The Effect of Tailored Information, Goal Setting, and Tailored Feedback on Household Energy Use, Energy-Related Behaviors, and Behavioral Antecedents." *Journal of Environmental Psychology* 27 (4): 265–76. DOI: https://doi.org/10.1016/j.jenvp.2007.08.002. [6]

Aigner, D. J., J. Newman, and A. Tishler. 1994. "The Response of Small and Medium-Size Business Customers to Time-of-Use (TOU) Electricity Rates in Israel." *Journal of Applied Econometrics* 9 (3): 283–304. DOI: https://doi.org/10.1002/jae.3950090305. [6]

Allcott, Hunt. 2011. "Rethinking Real-Time Electricity Pricing." *Resource and Energy Economics* 33 (4): 820–42. DOI: https://doi.org/10.1016/j.reseneeco.2011.06.003. [5, 6]

Allcott, Hunt, Benjamin B Lockwood, and Dmitry Taubinsky. 2019. "Regressive Sin Taxes, with an Application to the Optimal Soda Tax." *Quarterly Journal of Economics* 134 (3): 1557–626. DOI: 10.1093/qje/qjz017. [42]

Allcott, Hunt, and Sendhil Mullainathan. 2010. "Behavior and Energy Policy." *Science* 327 (5970): 1204–5. DOI: 10.1126/science.1180775. [7]

Allcott, Hunt, and Todd Rogers. 2014. "The Short-Run and Long-Run Effects of Behavioral Interventions: Experimental Evidence from Energy Conservation." *American Economic Review* 104 (10): 3003–37. DOI: 10.1257/aer.104.10.3003. [8]

Andor, Mark, Andreas Gerster, and Lorenz Goette. 2020. "Disaggregate Consumption Feedback and Energy Conservation." CEPR Discussion Paper DP14954. URL: https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3638059. [8, 53]

Attari, Shahzeen Z. 2014. "Perceptions of Water Use." *Proceedings of the National Academy of Sciences* 111 (14): 5129–34. DOI: 10.1073/pnas.1316402111. [7]

Attari, Shahzeen Z., Michael L. DeKay, Cliff I. Davidson, and Wändi Bruine de Bruin. 2010. "Public Perceptions of Energy Consumption and Savings." *Proceedings of the National Academy of Sciences* 107 (37): 16054–59. DOI: 10.1073/pnas.1001509107. [7]

Birzle-Harder, B., and K. Götz. 2001. *Grüner Strom: Eine Sozialwissenschaftliche Marktanalyse*. Studientexte des Instituts für Sozial-Ökologische Forschung. ISOE. URL: https://books.google.de/books?id=GC9wHAAACAAJ. [7]

Buchanan, Kathryn, Riccardo Russo, and Ben Anderson. 2015. "The Question of Energy Reduction: The Problem(s) with Feedback." *Energy Policy* 77: 89–96. DOI: https://doi.org/10.1016/j.enpol.2014.12.008. [36]

Chetty, Raj. 2009. "Sufficient Statistics for Welfare Analysis: A Bridge Between Structural and Reduced-Form Methods." *Annual Review of Economics* 1 (1): 451–88. DOI: 10.1146/annurev.economics.050708.142910. [9]

Chetty, Raj, Adam Looney, and Kory Kroft. 2009. "Salience and Taxation: Theory and Evidence." *American Economic Review* 99 (4): 1145–77. DOI: 10.1257/aer.99.4.1145. [7–9, 42]

Degen, Kathrin, Charles Efferson, Fabian Frei, Lorenz Goette, and Raffael Lalive. 2013. "Smart Metering, Beratung oder Sozialer Vergleich - Was beeinflusst den Elektrizitätsverbrauch?" Report for the Swiss Federal Office of Energy. University of Lausanne. [6, 15, 44]

DellaVigna, Stefano. 2009. "Psychology and Economics: Evidence from the Field." *Journal of Economic Literature* 47 (2): 315–72. DOI: 10.1257/jel.47.2.315. [8, 9]

Ehrhardt-Martinez, Karen. 2008. "Behavior, Energy, and Climate Change: Policy Directions, Program Innovations, and Research Paths." In. American Council for an Energy-Efficient Economy. [6]

Ehrhardt-Martinez, Karen, and John A. Laitner. 2010. "People-Centered Initiatives for Increasing Energy Savings." In. American Council for an Energy-Efficient Economy. [6]

Enke, Benjamin, and Thomas Graeber. 2019. "Cognitive Uncertainty." Working Paper 26518. National Bureau of Economic Research. DOI: 10.3386/w26518. [9]

Al-Faris, Abdul Razak F. 2002. "The Demand for Electricity in the GCC Countries." *Energy Policy* 30 (2): 117–24. DOI: https://doi.org/10.1016/S0301-4215(01)00064-7. [6]

Filippini, Massimo. 1995. "Electricity Demand by Time of Use. An Application of the Household AIDS Model." *Energy Economics* 17 (3): 197–204. DOI: https://doi.org/10.1016/0140-9883(95)00017-O. [6]

Finkelstein, Amy. 2009. "E-ztax: Tax Salience and Tax Rates." *Quarterly Journal of Economics* 124 (3): 969–1010. DOI: 10.1162/qjec.2009.124.3.969. [42]

Fischer, Corinna. 2008. "Feedback on Household Electricity Consumption: A Tool for Saving Energy?" *Energy Efficiency* 1 (1): 79–104. DOI: https://doi.org/10.1007/s12053-008-9009-7. [8, 36]

Gabaix, Xavier, and David Laibson. 2006. "Shrouded Attributes, Consumer Myopia, and Information Suppression in Competitive Markets." *Quarterly Journal of Economics* 121 (2): 505–40. DOI: 10.1162/qjec.2006.121.2.505. [6]

Hortaçsu, Ali, Seyed Ali Madanizadeh, and Steven L. Puller. 2017. "Power to Choose? An Analysis of Consumer Inertia in the Residential Electricity Market." *American Economic Journal: Economic Policy* 9 (4): 192–226. DOI: 10.1257/pol.20150235. [7]

Ito, Koichiro. 2014. "Do Consumers Respond to Marginal or Average Price? Evidence from Nonlinear Electricity Pricing." *American Economic Review* 104 (2): 537–63. DOI: 10.1257/aer.104.2.537. [5, 6]

Jessoe, Katrina, and David Rapson. 2014. "Knowledge Is (Less) Power: Experimental Evidence from Residential Energy Use." *American Economic Review* 104 (4): 1417–38. DOI: 10.1257/aer.104.4.1417. [6, 7]

Kőszegi, Botond, and Matthew Rabin. 2006. "A Model of Reference-Dependent Preferences." *Quarterly Journal of Economics* 121 (4): 1133–65. DOI: 10.1093/qje/121.4.1133. [38]

Laibson, David. 1997. "Golden Eggs and Hyperbolic Discounting." *Quarterly Journal of Economics* 112 (2): 443–78. DOI: 10.1162/003355397555253. [43]

Lijesen, Mark G. 2007. "The Real-Time Price Elasticity of Electricity." *Energy Economics* 29 (2): 249–58. DOI: https://doi.org/10.1016/j.eneco.2006.08.008. [6]

Patrick, Robert H, and Frank A Wolak. 2001. "Estimating the Customer-Level Demand for Electricity Under Real-Time Market Prices." Working Paper 8213. National Bureau of Economic Research. DOI: 10.3386/w8213. [6]

Reiss, Peter C., and Matthew W. White. 2005. "Household Electricity Demand, Revisited." *Review of Economic Studies* 72 (3): 853–83. DOI: 10.1111/0034-6527.00354. [5, 6]

Rutter, Richard, Konstantinos J Chalvatzis, Stuart Roper, and Fiona Lettice. 2018. "Branding Instead of Product Innovation: A Study on the Brand Personalities of the UK's Electricity Market." *European Management Review* 15 (2): 255–72. DOI: https://doi.org/10.1111/emre.12155. [7]

Steg, Linda, and Charles Vlek. 2009. "Encouraging Pro-Environmental Behaviour: An Integrative Review and Research Agenda." *Journal of Environmental Psychology* 29 (3): 309–17. DOI: https://doi.org/10.1016/j.jenvp.2008.10.004. [6, 43]

Sunstein, Cass R. 2015. "Behavioural Economics, Consumption and Environmental Protection," 313–27. DOI: https://doi.org/10.4337/9781783471270.00032. [6]

Taubinsky, Dmitry, and Alex Rees-Jones. 2017. "Attention Variation and Welfare: Theory and Evidence from a Tax Salience Experiment." *Review of Economic Studies* 85 (4): 2462–96. DOI: 10.1093/restud/rdx069. [42]

Tiefenbeck, Verena, Lorenz Goette, Kathrin Degen, Vojkan Tasic, Elgar Fleisch, Rafael Lalive, and Thorsten Staake. 2016. "Overcoming Salience Bias: How Real-Time Feedback Fosters Resource Conservation." *Management Science* 64 (3): 1458–76. DOI: https://doi.org/10.1287/mnsc.2016.2646. [9]

Ueno, Tsuyoshi, Fuminori Sano, Osamu Saeki, and Kiichiro Tsuji. 2006. "Effectiveness of an Energy-Consumption Information System on Energy Savings in Residential Houses Based on Monitored Data." *Applied Energy* 83 (2): 166–83. DOI: https://doi.org/10.1016/j.apenergy.2005.02.002. [6]

Wilson, Charlie, and Hadi Dowlatabadi. 2007. "Models of Decision Making and Residential Energy Use." *Annual Review of Environment and Resources* 32 (1): 169–203. DOI: 10.1146/annurev.energy.32.053006.141137. [6]

Chapter 2

Do Households Shift Electricity Consumption? Evidence from a Real-Time Pricing Tariff

Joint with Lorenz Goette

2.1 Introduction

Electricity suppliers and grid operators face a complicated logistical challenge: At all times, the grid-wide electricity demand has to be met. Electricity cannot easily be stored, and hence, has to be produced when it is needed. While most fossil fuels can be used when needed, most renewable energy sources, such as solar and wind energy, are not available at all times. In addition, their fluctuating availability cannot be fully predicted. In the power market, predictable fluctuations in electricity demand and supply are usually offset by fuel plants producing electricity if needed. Short-time fluctuations are met by oil and gas plants, which can be started and stopped relatively quickly. Using such back-up plants to offset fluctuations in energy availability and demand, however, requires additional capacities in electricity production that are rarely used. Their marginal (and marginal social) costs of production, however, are high, especially at full capacity. So far, policymakers have mostly been focusing on adapting the electricity supply to the electricity demand. However, adapting electricity demand to the electricity supply could ease the problem and would reduce the need for back-up plants, increasing the efficiency in the electricity market.

In this paper, we examine the effectiveness of residential real-time electricity pricing (RTP) as a mechanism to adjust electricity demand to the contemporaneous electricity supply. In RTP schemes, households do not pay a fixed price for electricity at all times. Instead, electricity prices vary by hour in non-fixed patterns - for example, to follow the marginal costs of electricity provision, which vary in real-time.[1] We investigate the RTP mechanism set by an electricity provider operating in Austria and Germany, which passes wholesale electricity prices from the European Power Exchange (EPEX) on to the customer households. These prices are generated based on supply and demand predictions and are set for each hour of the next day. Consequently, households in our sample do not pay a flat rate, but hourly changing electricity prices that, at each hour, reflect the marginal costs of electricity provision.

Time-varying energy pricing (TVP) has often been proposed to close the gap between the marginal costs of energy provision and the electricity price (Ito, Ida, and Tanaka (2018)). Yet, TVP mechanisms are still uncommon (Borenstein (2005)). Customers are used to being charged a flat rate, and they are naturally suspicious of new pricing schemes. Hence, while electricity providers and grid operators are generally aware of the beneficial impact of dynamic pricing on consumption smoothing over the day, they often fear the potential backlash of customers if the price design is perceived as too complicated or too opaque. On the other hand, the idea of dynamic electricity pricing is not new. The daytime-/nighttime-tariffs used by many electricity providers are motivated by the same logic: With total electricity consumption dropping during the night, plants that cannot easily adjust their capacity over the day overproduce. In order to provide incentives to use more electricity during the night and less during the day, electricity providers charge less per kWh during the nighttime hours.[2] Introducing customers to dynamic pricing mechanisms with time-varying prices can provide incentives to reduce or shift consumption from peak-hours to off-peak hours, smoothing total electricity consumption over the day. In contrast to most TVP mechanisms, the RTP mechanism in our setting generates electricity prices that change hourly in non-fixed patterns, but are known since the day before.

We use individual-level hourly consumption data of almost 900 German households, which were exposed to electricity RTP based on EPEX wholesale prices. Our observation window spans the period between April 1, 2019 and December 31, 2020. In our sample, we identify four clusters of households that can be categorized with respect to their electricity consumption profile over the day. Excluding households with extremely high overall and large nighttime consumption leaves us with our final sample of 830 households.

As the wholesale electricity price is the result of grid-wide electricity production and demand, it is endogenous with respect to domestic consumption patterns. To address this challenge, we instrument for the hourly electricity price using hourly wind production in Germany. Hourly wind energy production exhibits a strong neg-

1. Customers usually know the electricity prices they are facing in advance.

2. Such daytime- and nighttime tariffs are called time-of-use (TOU) pricing designs. These mechanisms feature electricity prices that are higher during peak periods and lower during off-peak periods, but (in contrast to RTP) move according to a predictable pattern. Peak periods can be defined as seasons, months or - in the case of daytime-/nighttime tariffs - time of the day.

ative relation with the hourly electricity price on the first stage: One additional GW in wind production is associated with a 0.1 cent-drop in the electricity price. On the other hand, wind energy production is unlikely to be correlated with domestic electricity consumption patterns, conditional on the electricity price (especially when controlling for local weather). In addition to our analyses, we run the same analyses on a panel of 130 control households spanning the period between January 1, 2019 and January 31, 2020.

We find that households in our sample, in contrast to households in our control sample, show a strong overall sensitivity to hourly changing electricity prices, with an estimated overall price elasticity of -0.676. Estimated price sensitivities are stronger for low electricity prices. More importantly, we find that households show stronger demand reactions in response to the electricity price varying over the day than in response to the electricity price varying across several days: Zooming in on our results by adding date fixed effects, we estimate an intra-day price elasticity of -1.668. By adding date fixed effects, we only exploit variation in the electricity price that occurs within the day. Zooming out of our results by collapsing electricity consumption on a daily level (and thus ignoring any within-day price variation) delivers an estimate for the inter-day price elasticity of -0.453, indicating that the scope for demand responses to price variation crucially depends on the time horizon of consumption. Households in our sample react more strongly to price variation over the day than price variation across days by shifting electricity consumption within the day rather than over several days. Possible examples for this could be showering in the morning instead of showering during the expensive evening hours or cooking earlier. The scope to shift such activities over several days is limited. These results are robust to a large number of different specifications and sample restrictions and, although the estimated price sensitivities are larger on weekends, are still highly significant during weekdays. Finally, estimating household level price sensitivities reveals that our estimates are driven by a large share of households, rather than just a few highly elastic ones.

As households in our sample self-selected into the real-time pricing scheme, we also assess the external validity of our results. Since we have no household-level information about the households in our sample other than their zipcode, we use information on sociodemographic and socioeconomic variables on the zipcode level, provided by the data provider *microm* to investigate their sociodemographic and socioeconomic environment. While we find statistically significant differences between Germany and our sample (for example, in the share of homeowners, education, income and family structures), we also find that only for a few of these variables the differences are large in absolute terms. Still, we report that households in our sample live in zipcode areas with a higher average socioeconomic status and more children-centered family structures.

Our findings provide real-time demand elasticity estimates for different time horizons, thus contributing to the evaluation of the effectiveness of different types of dynamic electricity pricing. In order to estimate the effectiveness of such pricing mechanisms and develop an appropriate pricing system, it is crucial to know the consumers' demand responsiveness with respect to the electricity price (Reiss and White (2005)). However, universally applicable empirical work on price elasticities is scarce. While there is a large strand of empirical literature on electricity price elasticities, different data structures yield different estimates for the price sensitivity. Andruszkiewicz, Lorenc, and Weychan (2019) provide an elaborate overview of estimated price elasticities of electricity demand analyses. Most studies analyze aggregate electricity consumption on an annual or quarterly level (Holtedahl and Joutz (2004), Boonekamp (2007), Alberini and Filippini (2011), Boogen, Datta, and Filippini (2017), Filippini (2011), Okajima and Okajima (2013)), with the estimated price elasticities depending on the granularity of the available data.

Our study, which uses individual-level hourly consumption data, contributes to the literature on household price elasticities (Harding and Sexton (2017)). Only few studies use individual-level consumption data (Schulte and Heindl (2017), Silva, Soares, and Pinho (2018), Boogen, Datta, and Filippini (2014), Volland and Tilov (2018)) to estimate price elasticities, and even fewer use hourly consumption data (Lijesen (2007), Patrick and Wolak (2001), Allcott (2011)).

Other studies analyze individual-level, high-frequency consumption data, but instead of RTP rather focus on critical peak pricing (CPP) mechanisms, which introduce (often infrequent, unanticipated) price spikes (Bollinger and Hartmann (2020), Faruqui and George (2005), Faruqui, Sergici, and Akaba (2012), Jessoe and Rapson (2014), Ito, Ida, and Tanaka (2018)). Additionally, since most of the previous empirical work on high-frequency demand responses stems from field experiments, previous estimates are naturally limited with respect to sample sizes and thus, precision.

Two notable studies that explicitly address both household-level, residential electricity consumption and hourly changing prices are Allcott (2011) and Fabra, Rapson, Reguant, and Wang (2021). Allcott (2011) analyzes an energy pricing plan in Chicago in 2003, which exposes residential consumers to RTP. Using hourly household-level consumption data, he finds significantly negative price elasticities, which, however, stem from energy conservation during peak price hours, with no net increase in consumption when the price is low - indicating a limited scope for shifting consumption over the day. He also emphasizes the importance of technology facilitating the access to hourly prices to make RTP a viable option. Fabra et al. (2021) are among the first to analyze a large-scale roll-out of RTP. Starting in 2015, households in Spain defaulted into an RTP tariff, setting the hourly electricity price according to the wholesale electricity market outcomes. They find an average price elasticity of zero, listing lack of awareness of RTP, costly information acquisition, and small potential gains of demand response to price variation as possible reasons. In contrast to their sample, households in our sample actively sought out a real-time pricing scheme and are thus less likely to be constrained by these barriers.

The remainder of this paper is organized as follows: In Section 2.2, we explain in detail the process at the European Power Exchange EPEX, in which electricity wholesale prices are determined. Section 2.3 gives an overview over the data used. Section 2.4 presents descriptive statistics and discusses the representativeness of our sample. Section 2.5 lays out the empirical challenges of the estimation framework discusses our empirical approach. In Section 2.6, we present the results. Section 2.7 concludes.

2.2 The European Power Exchange

To get a better grasp of our empirical strategy, it helps to take a closer look at how electricity prices are generated. In Germany, approximately 75% of the electricity price consist of flat-rate taxes and fees, which are fixed over time. The remaining 25% consist of the wholesale price, at which electricity is sold at the European Power Exchange *EPEX*.

At the current stage, electricity cannot be easily stored at a large scale. It has to be used when it is produced. This is why electricity production usually has to move with electricity demand to avoid power outages. As the energy grid does not allow for large deviations in power load, the largest entities in the energy market (the largest electricity producers and consumers) announce, how much energy they are going to produce or consume at each time of the next day. From these reports the grid operators can roughly calculate the electricity load of their grids for the next day. This process is called *Dispatching*. If this process reveals large, systematic disparities between electricity production and demand at any point, these can then be easily balanced out. Small short-run deviations in electricity consumption and production due to unforeseen consumption spikes or weather anomalies are balanced out in a process called *Redispatching*.

The electricity itself is traded at the EPEX. The European Power Exchange is the electric power exchange for central Europe, including Austria, Belgium, Denmark, Finland, France, Germany, Great Britain, Luxembourg, the Netherlands, Norway, Poland, Sweden and Switzerland. Here, electricity is traded via three different channels, each targeting a different time horizon:

- *Power Derivatives Market*: In the Power Derivatives Market long-run energy deliveries up to six years in the future are traded.
- *Day-Ahead-Market*: Day-ahead markets are operated through a blind auction which takes place once a day for all hours for the following day, during each day of the year. Before the order book closes at 12 p.m., market participants log two types of orders into the system: Firstly, orders for each time interval for the next day reflecting their willingness to buy or sell as a volume for all relevant price ticks. Secondly, block orders spanning multiple time intervals. The algorithm employed by the EPEX then generates a demand curve based on the buy-

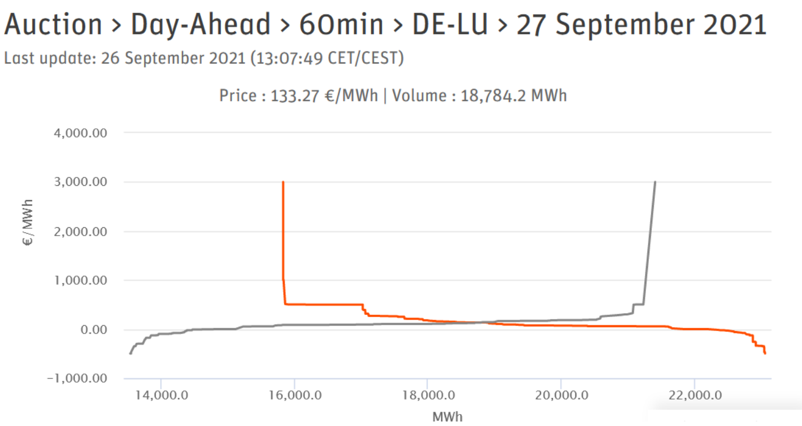

Figure 2.1. Supply and Demand Curves at EPEX

This figure shows a screenshot from the EPEX-website, displaying the supply and demand curves generated at the EPEX on Sep 26, 2021, and the resulting market clearing price for the period from 1 p.m. to 2 p.m. Source: EPEX (https://www.epexspot.com/en/market-data?market_area=DE-LU&trading_date=2021-12-18&delivery_date=2021-12-19&underlying_year=&modality=Auction&sub_modality=DayAhead&product=60&data_mode=aggregated&period=, last visited Sept 27, 2021)

orders and a supply curve based on the sell-orders for each hour of the following day. The market clearing prices (MCPs) are then defined at the intersections of both curves for each hour of the day (See Figure 2.1). These prices are available to households in our sample at 2 p.m. for the next day. Once the MCPs are determined, there is one price for each market participant, i.e. all buyers who submitted volumes at a higher price pay only the MCP and all sellers who submitted orders at a lower price still receive the MCP. Consequently, like in most textbook commodity markets, the electricity price at the Day-Ahead-Market is determined by energy production and energy demand.

- *Intraday-Market*: The Intraday-Market is the most fine-grained form of electricity trading: Market participants trade continuously, while electricity is delivered on the same day. Trades are executed as soon as buy- and sell-orders match, and can be executed up to 5 minutes before delivery. The Intraday-Market is usually used to balance out sudden deviations in energy demand and is an essential tool in grid load control.

In our analysis, we will focus on the day-ahead market prices, which are used by the electricity supplier and are thus fixed and available at the electricity provider's website at 2 p.m. the day before delivery.

2.3 Data

We use hourly data provided by an electricity provider operating in Germany and Austria. According to the company website, the electricity provider is one of the first providers in Germany with hourly moving electricity prices. In order to ensure

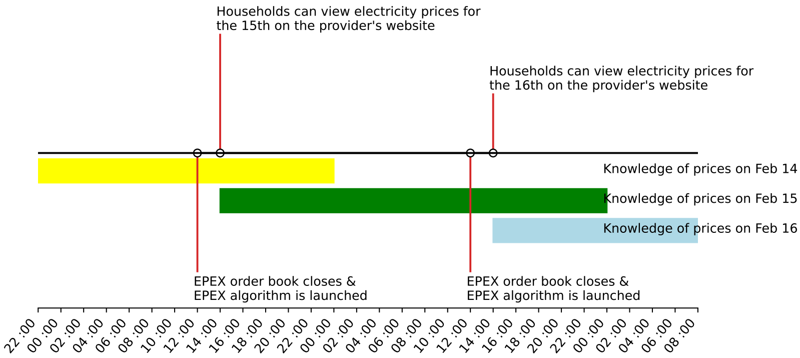

Figure 2.2. Information Set Timeline

The figure shows which information sets are available to the households at every time. Our example time-frame covers February 13, 10 p.m. to February 16, 8 a.m.
At 12 p.m., the EPEX order book closes and the algorithm, which generates the demand- and supply curve, is launched and produces 24 market clearing prices for the next day. These prices are made public to the households on the electricity provider's website at 2 p.m. Hence, from 2 p.m. onwards, a household knows all electricity prices for the following day. The electricity prices for the current day are already known since the previous day, 2.p.m.

a maximum level of transparency in the price setting process the electricity provider sets the market clearing prices generated at the EPEX as net prices for the customers. With taxes and fees, the variable share of the electricity price is approximately 25%. This pricing design is captured in the HOURLY-tariff. In order to be eligible to use this tariff, customers need to have a smart meter installed to capture high-frequency electricity consumption. This ensures that the provider can assign the customers' consumption to each hour of the day. Customers can see the electricity prices for the next day at 2 p.m. on the provider's website. A timeline describing when households know which prices is displayed in Figure 2.2.

As the electricity provider is operating in Austria and Germany, we have data on customers from both countries. Our data include each household's electricity consumption in 15-minute intervals between April 1, 2019 and December 31, 2020 as well as the hourly day-ahead prices determined at the EPEX. The panel consists of approximately 900 households in Germany and 1'200 households in Austria. As the resulting hourly panels over such a long period are very large, we will primarily focus on the German dataset.

We then merge hourly electricity production data provided by AGORA[3] to our dataset and use the households' zipcode to merge daily local weather data[4], pro-

3. AGORA Energiewende (2021). The original data can be requested at info@agora-energiewende.de See also https://www.agora-energiewende.de/service/agorameter/chart/power_generation/01.01.2019/31.12.2020/, last visited Jan 13, 2022

4. For this, we use information on the five weather stations closest to a household's zipcode maintained by the German Meteorological Service: https://opendata.dwd.de/climate_environment/CDC/observations_germany/climate/daily/kl/historical/, last visited July 24, 2021

vided by the German Meteorological Service and sociodemographic information at the zipcode level, provided by the data provider *microm*, to our dataset. Thus, while we do not have household-level information about the households in our sample, we have information about the socioeconomic environment each household lives in. This information contains, for example, the share of homeowners, electric cars and non-German household heads as well as age, family, education and income structures at the zipcode level. The Austrian dataset additionally contains information on whether a household used a solar panel or a heatpump. In total, our panel contains consumption data for 1'204 households in Austria and 899 households in Germany between April 1, 2019 and December 31, 2020 totaling approximately 6 million observations in Germany and 13 million observations in Austria. From now on, unless mentioned otherwise, we restrict our analysis on the German sample for computational ease.

Finally, we use a panel of 130 households, spanning the period between January 1, 2019 to January 31, 2020, from a field study conducted in Germany.[5] The households in this sample did not receive any treatment in the original study, but were equipped with smart meters to serve as a control group. In particular, they were not exposed to real-time pricing and hence, by providing high-frequency consumption data without being exposed to treatment, compose the control group in our empirical approach. Because, for households in this sample, we have geolocation data at the state level, we use the households' state to merge daily weather data to the control sample.

2.4 Descriptive Statistics

Before diving into the analysis, we take a closer look and compare the households in our sample to the average German household. Unfortunately, we have no individual-level information on household characteristics other than the zipcodes. We can, however, analyze where the households in our sample live to rule out geographical selection and use aggregate socioeconomic variables on the zipcode level to inform about the socioeconomic environment our sample households live in. Additionally, we can compare electricity consumption in the analyzed households to average consumption profiles.

2.4.1 Hourly Consumption Profiles

Figure 2.3 displays the average electricity load profile over the day for the German households in our sample, the average electricity load profile for control households in a field experiment in Zurich, Switzerland in 2012 (Degen, Efferson, Frei, Goette,

5. doi:10.5281/zenodo.3855575. See Beyertt, Verwiebe, Seim, Milojkovic, and Müller-Kirchenbauer (2020) for the whole paper.

and Lalive (2013)), two standard consumption profiles for German households provided by an electricity provider located in Berlin, Germany and the average load profile of the 130-household control group of a study conducted in Germany[6]. 95% confidence intervals are also included. Both German standard consumption load profiles were generated based on standard consumption data provided by a Berlin-based electricity provider.[7] The standard consumption data was then scaled down to fit the average annual electricity consumption of a two-person- and a three-person household, respectively.[8] Note that the generated standard consumption profiles do not distinguish between households of different sizes other than through the magnitude of consumption. That is, consumption patterns are assumed to be the same across households for the standard profiles.

The German standard profiles clearly reflect familiar daily patterns in households: People get up between 5 and 8 a.m. and start preparing the day: With showering and preparing breakfast electricity consumption jumps up quickly in the morning. At noon, the first spike is reached, when children are returning from school and lunch is prepared. Consumption then first decreases until people return from work, start preparing dinner etc. This results in a second consumption spike between 6 and 10 p.m. This standard load profile is fairly consistent over time, with differences arising from different daily schedules on weekends and higher average consumption during the winter months, but the overall bimodal shape with a stronger spike in the evening hours is common in most industrial nations.[9] We can also see that the load profiles generated by the German and Swiss control households closely resemble the consumption patterns suggested by the German standard profiles, though the absolute consumption magnitude differs.

When comparing the load profile of the average German or the average Zurcher household with the average load profile of the households in our sample, two results are immediate: Firstly, we can observe that overall consumption in our sample is much higher than for the remaining households, even if we assume an average household size of three people.[10] Additionally, the shapes of the load profiles differ.

6. doi:10.5281/zenodo.3855575

7. https://www.stromnetz.berlin/netz-nutzen/netznutzer, last visited Sep 27, 2021

8. The average annual electricity consumption of a two-person household is 2'500 kWhs, that of a three-person household 3'500 kWhs. The average household size in Germany is 1.99. Source: Bundeszentrale für Politische Bildung, https://www.bpb.de/nachschlagen/zahlen-und-fakten/soziale-situation-in-deutschland/61584/bevoelkerung-und-haushalte, last visited Jan 5, 2022

9. Still, electricity providers use a polynomial to generate the data on which these graphs are based. That is, the underlying consumption pattern over the day is adapted with respect to days of the week, months of the year and cyclical patterns. This dynamization generates the profiles' standard errors.

10. Note that it is plausible to assume a higher-than-average number of household members, as owning a smart meter, a prerequisite to use the HOURLY-tariff, is more frequent among homeowners than among tenants. This can explain the higher overall consumption in our sample compared to the control households. Additionally, as shown in Table 2.A.4, households in our sample live in zipcode areas with more family- and children-centered family structures.

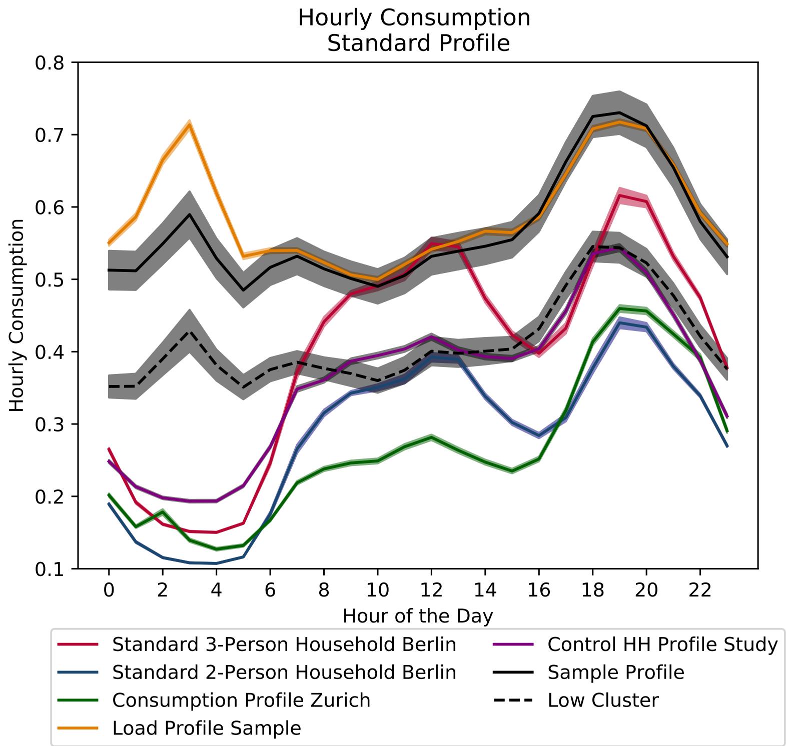

Figure 2.3. Daily Consumption Profile: Sample vs. Average German Household

This figure displays different average electricity load profiles over the day (in kWh per hour). That is, it shows average consumption for each hour of the day. The orange line represents the load profile for all households in our sample. The black solid line represents households in our regression sample. For our analyses, we excluded households with extremely high consumption and an exceptional spike in nighttime consumption by using a K-Means clustering algorithm based on a household's daily consumption profile. The black dashed line represents households in the "low consumption" cluster of our sample. The blue and red lines represent a standard two- and three-person household, respectively, according to the German energy provider. Finally, the green and purple profiles represent control households from smart metering studies in Zurich, Switzerland (Degen et al. (2013)) and Germany (Beyertt et al. (2020)), respectively. The latter comprise a control sample for our study. The light bands denote 95% confidence intervals.

While the daytime- and evening-consumption patterns are somewhat similar, we observe a remarkable nighttime consumption spike in our sample. Since our sample does not contain households with the regular fixed electricity price, we cannot easily tell whether the differences in the consumption profiles are due to differences in the composition of the households in our sample and the average German household or due to the pricing mechanism. In order to find out if the household composition indeed differs from the average German household composition in terms of consumption patterns, we generate clusters of households in Section 2.5.1, based on consumption patterns to uncover if the unusual load profiles are driven by just a few households. Indeed, we find that the majority of the households in our sample exhibit load profiles similar to the control households. The "low-consumption"-cluster, represented by the dashed line in Figure 2.3, which makes up more than 67% of

Table 2.1. Hourly Electricity Consumption

	Total Sample[a]	Two Person	Three Person	Ext. Study	Zurich	Sample[b]
Mean Consumption	.591	.285	.399	.360	.265	.530
	(.000)	(.001)	(.001)	(.000)	(.000)	(.000)
Median Consumption	.308	0.298	.417	.221	.130	.300
25-Percentile	.125	0.179	.250	.117	.070	.123
75-Percentile	.645	0.363	.508	.427	.300	.612
Observations	6’004’525	17’543[c]	17’543[c]	1’235’390	1’057’891	5’583’500
Households	899	1[c]	1[c]	130	1’009	830

[a]Including all Households
[b]Including only Households in the Regression Sample
[c]Note that this is a sample profile based on representative consumption data.

Notes: This table displays average consumption and 25-, 50-, and 75-percentiles of hourly household consumption from different data sources. In the first column, we show consumption in our total sample, columns 2 and 3 show consumption from a representative household profile over the sample period, based on consumption data from a Berlin-based electricity provider. Column 4 shows consumption in our control sample, consisting of 130 households from a different study conducted between January 1, 2019 and January 31, 2020 in Germany. Column 5 shows consumption for control households in a field experiment in Zurich, Switzerland in 2012 (Degen et al. (2013)). Column 6 shows consumption in our final regression sample, which excludes households with extremely high overall and high nighttime consumption.

the households in our total sample, still exhibits a remarkable nighttime consumption spike, but total consumption is much lower. In our final analysis, we will focus on the two clusters with the lowest consumption in our sample, represented by the solid black line, in order to ensure the representativeness of our sample in terms of baseline consumption.

2.4.2 Hourly Price Profiles

We will also analyze the day-ahead electricity prices generated at the EPEX. As outlined in Section 2.2, the day-ahead electricity price is a result of forecast electricity production and demand. Figure 2.4 shows that this results in an hourly electricity price profile that closely resembles the electricity profile of the average German household. As we would expect from the pricing mechanism, the electricity price is high when overall consumption is high and low when overall consumption is low. As the electricity price is also influenced by industrial electricity demand, the first peak in the electricity price is reached earlier than the household consumption peak, namely, when industrial electricity consumption hits the high level it maintains until the evening. The second spike is then mostly caused by the spike in household consumption in the evening.

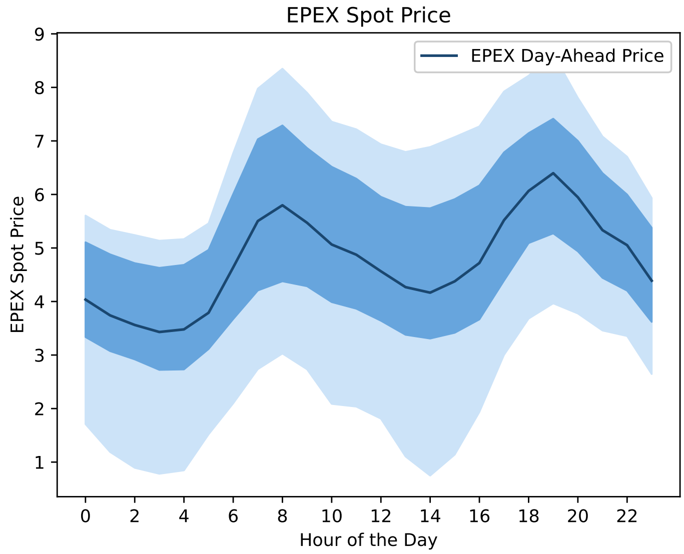

Figure 2.4. Electricity Price Profile

This figure shows the average net EPEX electricity price profile over the day between January 1, 2019 and December 31, 2020. The black line denotes average EPEX prices over the day in ct/kWh. The blue bands denote the 10, 25, 75, and 90-percentiles of EPEX prices.

2.4.3 Representativeness of Households in the Sample

The high average consumption and the shape of the consumption load profile of the households in our sample may raise questions about the representativeness of our sample. Especially the high consumption magnitude of households in our sample appears to limit the external validity of our analysis. Two remarks are in order: In order to be eligible for the HOURLY-tariff, households need to have an installed smart meter - otherwise the electricity provider would not be able to bill households according to their hourly consumption. So far, the only consumers who are by law required to have a smart meter installed (regardless of their energy provider), are households using more than 6000 kWhs per year, producer-consumers producing more than 7 kW using, for example, solar panels, and households with an interruptible meter.

However, as representatives of one of the largest smart meter-suppliers in Germany confirm, the share of homeowners (and thus, households with more-than-average members) among households with a smart meter is very high. This has several reasons:

- Tenants are still hesitant to install a smart meter if it is not required, as many believe that the electricity meters of each apartment belong to the landlord. A frequently asked question for smart meter-providers is whether their landlord can prohibit the installation (which they cannot).

- Landlords, on the other hand, have no incentive to proactively install a smart meter, as the electricity bill is usually paid by the tenant.
- Smart Meters cannot easily move if the tenant does. Hence, if a tenant moves to a new apartment, they have to install a new smart meter, if it is not installed, yet.

Hence, with a higher share of homeowners among smart meter users, we expect the share of homeowners in our sample to be very high. Based on consumption data from Zurich between 2011 and 2013, we can show that homeowners exhibit much higher electricity consumption than tenants.[11] As Table 2.1 shows, average consumption in Zurich was roughly .265 kWhs per hour. For homeowners, average consumption was much higher, with .331 kWhs, which is almost 25% higher than overall average consumption, and more than 40% higher than average consumption for tenants (.231 kWhs per hour). We thus believe that the high electricity consumption in our sample is most likely be driven by an overrepresentation of homeowners in our sample.

As the electricity provider's customers are self-selecting into our sample, we need to make sure they do not represent a special kind of niche customers. Unfortunately, we only have information regarding the households' zipcodes, with no individual-level demographic information. However, as Figure 2.5 shows, households in our sample are not restricted to a geographical area, although the share of customers is slightly higher in western Germany than in eastern Germany.

Based on the households' zipcodes, we use information on sociodemographics on the zipcode level, obtained from the data provider *microm*[12] to characterize the socioeconomic environment of the households in our sample. Based on this information, we analyze how representative households in our sample are. In Appendix 2.A, we analyze the representativeness with respect to sociodemographics, education, income, family structures and socioeconomic status. As Table 2.A.1 shows, zipcodes with a high share of homeowners are significantly overrepresented in our sample, as are zipcodes with a high share of households, in which the household head has reached the A-Level (*Abitur*). Conversely, zipcodes with a high share of non-German household heads and zipcodes with a high share of household heads over the age of 60 or under the age of 30 are underrepresented. The zipcodes in our sample do not significantly differ with respect to the share of electric or hybrid cars. Table 2.A.2 reveals that the zipcodes in our sample also slightly differ with respect to income. Overall, the zipcodes in our sample seem to exhibit a higher average income than the German average. Tables 2.A.3 and 2.A.4 show that in the zipcodes in our sample the share of families is significantly higher than the German average. In line with

11. This is most likely due to the higher number of household members.

12. microm Micromarketing-Systeme und Consult GmbH. The datasets used here can be requested under info@microm.de

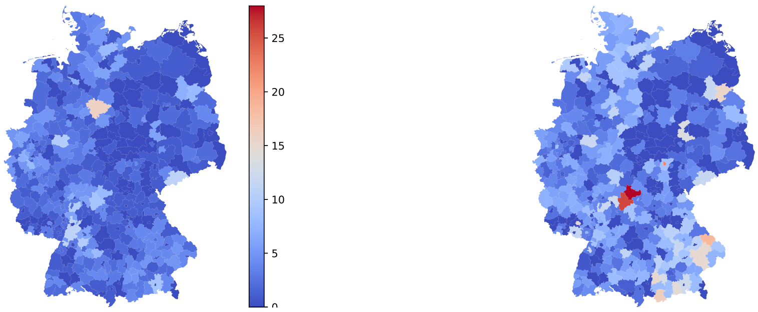

(a) Heatmap of Household Locations

(b) Heatmap of Household Locations Relative to Population

Figure 2.5. Wind Production and Energy Prices

This figure shows a heatmap of where the households in our sample are located on a county level, based on their zipcode information. Panel (a) is based on the absolute number of households in our sample per county. Panel (b) is based on the number of sample households in each county relative the number of residents in that county.

the previous findings, we find in Table 2.A.5 that households in the zipcodes in our sample fall into higher socioeconomic categories than the German average. Overall, the results differ significantly between Germany and the zipcodes in our sample, but we also see that households in our sample are not limited to only a few niche zipcodes in Germany.

In conclusion, we find significant differences in sociodemographics, education, income- and family structures between the German average and the zipcodes in our sample. As expected, households in our sample live in zipcode areas with a higher share of homeowners and families, higher average education, income and socioeconomic status, a lower share of non-German household heads and households heads above the age of 60 or below the age of 30. We thus cannot say that our sample is representative of German households. However, also note that, while statistically significant, the absolute differences are relatively small for most variables analyzed. The only differences we describe as economically significant are the differences in the share of homeowners, income and family structures.

Finally, note that to draw meaningful inference on the policy implications of RTP, households in our sample do not necessarily have to be representative of the German population. While the external validity of our sample cannot be guaranteed, the policy implications are not automatically affected by this. With the planned roll-out of smart meters in every household in Europe[13], it is possible to offer the option for RTP to every household and let them self-select into treatment. According to a survey conducted by Forsa, only 40 percent of consumers have ever heard of time-varying electricity prices, but more than half are willing to use TVP tariffs. Almost

13. Source: European Commission, https://ses.jrc.ec.europa.eu/smart-metering-deployment-european-union, last visited Jan 6, 2022

30 percent expect to financially benefit from TVP, and more than 60 percent would be willing to shift electricity consumption from peak consumption hours to off-peak hours in response to TVP. Young consumers are especially interested in such tariff models (Forsa (2015)), indicating that acceptance could be high if such tariffs were offered to customers.

2.5 Empirical Strategy and Results

2.5.1 Clustering Households

In order to identify the households that drive the surprising shape of the electricity load pattern shown in Figure 2.3, we generate clusters of households according to electricity consumption during different times of the day. We calculate average hourly electricity consumption for each household and use K-means clustering, one of the most popular unsupervised machine learning algorithms, to partition households into clusters of households with similar consumption patterns over the day. K-means clustering is an algorithm used to partition n data points into a fixed number of clusters, while trying to minimize the residual sum of squares within the clusters (*Within Sum of Squares WSS*), that is, trying to minimize the following expression:

$$WSS = \sum_{k=1}^{K} \sum_{x_i \in c_k} ||x_i - \mu_k||^2 \tag{*}$$

Where μ_k denotes the cluster centroid and x_i denotes individual i's data point. In our case, x_i is a 24-dimensional vector of electricity consumption during each hour of the day. The algorithm then proceeds as follows:

- 1. *Randomly* choose K households from the set of households with average consumption vectors $x_1, \cdots, x_K$
- 2. Assign each household i with average consumption vector x^i to the cluster, for which the second sum of expression (*) is minimized. The initial cluster centroid is just the consumption vector of the initial household in the cluster from step 1.
- 3. Calculate the new cluster centroid, including the newly assigned households
- 4. Repeat steps 2 and 3 until the assignment of households to clusters does not change anymore

The K-means clustering algorithm is a very appealing approach to handling the heterogeneity in household's consumption patterns (Trotta (2020)). Its simplicity and ease of implementation ensure transparency and allow for analyses of large datasets. However, despite falling into the category of unsupervised learning algorithms, the number of clusters K used for the K-means algorithm has to be chosen by the researcher. We choose to generate four clusters of households. This approach has been used by Trotta (2020) after applying the same argument to danish domestic electricity consumption data.[14]

Using the clustering algorithm, we can classify households according to four different consumption profiles.

- Households exhibiting an average consumption profile similar to the consumption profile of the average German household, both in terms of consumption magnitude and pattern. They make up more than 67% of the households in our sample. (Low Cluster)
- Households exhibiting a similar consumption pattern, but with higher absolute magnitude. We will label these medium-consumption households, though their consumption magnitude is significantly higher than consumption for the control households. They make up almost 25% of the households in our sample. (Medium Cluster)
- Very high consumption households with consumption 5 to 6 times as high as the average German household. These households also exhibit a different load pattern over the day, with consumption staying relatively low over the night, increasing in the morning and staying constant until the evening. They make up 6% of the households in our sample. (High Cluster)
- Households using more energy than the average German household, with a remarkable spike during the night (Nighttime Cluster). These households mainly drive the surprising consumption spike during the night depicted in Figure 2.3, despite making up less than 1.5% of the households in our sample. (Night Cluster)

Figure 2.6 displays the average electricity load profiles for each of the four clusters of households generated by the clustering algorithm. The high nighttime consumption for the nighttime cluster is most likely driven by the usage of interruptible electricity meters (*IM*) or an Application Programming Interface (*API*). These allow households to automatically turn on electric devices or charge the electric car when the price is low, which is usually during the night. The electricity provider offers an option to automatically exploit the energy price variations. On their website, they host an API-datafeed providing the electricity prices for the next day, which can then be fed to smart household appliances. Additionally, the provider collaborates with several suppliers of heatpumps, that are able to recognize the electricity prices.

Unfortunately, the German data set does contain information on whether a household uses an electricity meter with an API or whether it uses a heatpump. Luckily, we have information on the usage of interruptible electricity meters (IMs)

14. Note that in our case, including more clusters only adds more clusters with load profiles between the low- and medium cluster load profiles, which does not add much meaningful information.

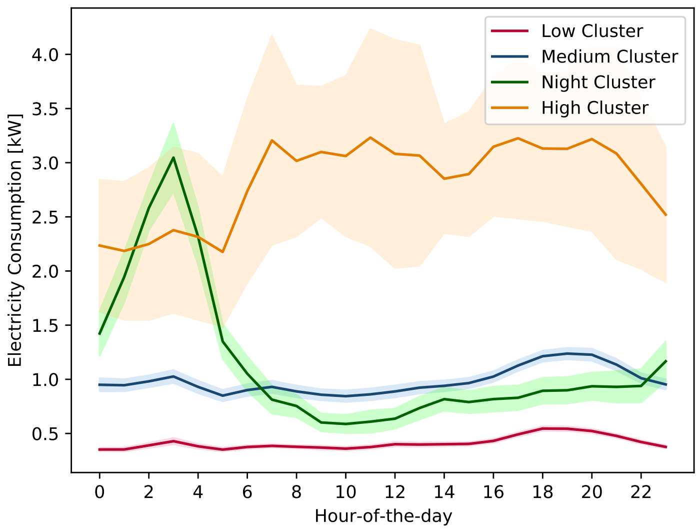

Figure 2.6. Daily Consumption Profiles by Clusters

This figure shows the average electricity load profiles of household clusters generated by the K-Means clustering algorithm. The high- and nighttime cluster will be excluded from further analyses, unless otherwise specified. Analyzing the high-cluster households yields limited external validity due to the extremely high baseline consumption, while the spike in nighttime consumption for the nighttime cluster households is likely driven by heatpumps. The light bands denote 95% confidence intervals.

and heatpumps in the Austrian data. We have information on whether a household uses an IM, which in turn is usually linked to an electronic heatpump. Since heating is very energy intensive, the spikes from using electricity through the IM are usually clearly visible in the load profiles of households with IM. Figure 2.7 displays the electricity consumption profiles from the Austrian dataset for households with and without an IM. We clearly see the pronounced consumption spike during the night for households with an IM, while consumption over the day is substantially lower than for households without IM. This indicates that excluding the nighttime cluster from Figure 2.6 from our sample indeed helps to alleviate concerns regarding the role of automated demand responses. While we do not have information whether households in our sample use a heatpump, the similarity to the load profile of households in the Austrian sample makes us confident that excluding households from the nighttime cluster excludes any automated demand response.

Since the use of IMs and APIs does not constitute a behavioral reaction to real-time electricity pricing, we will exclude households from the nighttime cluster from the analysis. 874 households remain. Additionally, we exclude households from the "very high"-consumption cluster, as they are extreme outliers, even in a sample with very high average consumption, giving us a final RTP-sample of 830 households. However, in Appendix 2.D.4, we also conduct our main analyses on the full sam-

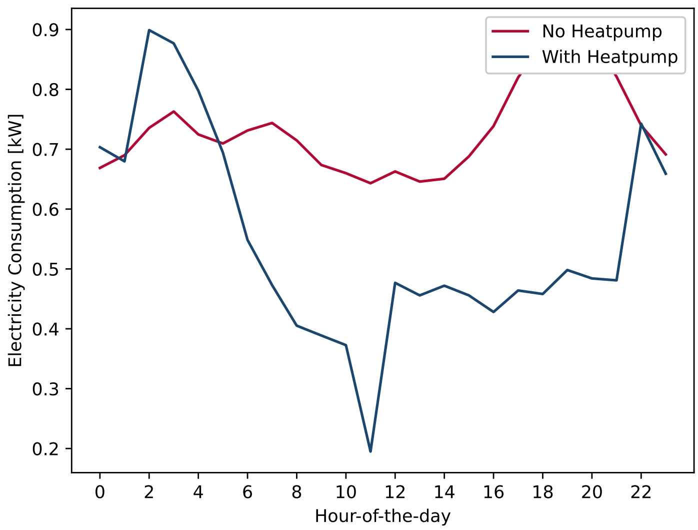

Figure 2.7. Daily Consumption Profiles with and without Heatpump

This figure shows the average electricity load profiles for households in Austria with and without heatpump.

ple and in Appendix 2.D.5 we do the same using only households from the "low consumption" cluster.

2.5.2 The IV Approach

Since the day-ahead market clearing prices (MCPs) for each hour-of-the-day are defined at the intersections of the respective supply- and demand curves, a simple regression of hourly electricity consumption on the real-time electricity price would yield biased estimates of price reactions. As shown in Figure 2.4, prices are high when overall consumption is high, that is, in the evening and especially during the cold season. Even though we don't expect the small number of households in our sample to have any significant effect on grid loads and thus electricity prices, their consumption patterns still resemble those of most households. Since electricity consumption patterns over the day influence the electricity price, we face a classical problem of reversed causality.

The empirical challenge inherent to this is to find a suitable instrument for the electricity price, which influences electricity consumption only through its effect on the electricity price, but is otherwise uncorrelated with domestic electricity consumption. We focus on the supply side of the electricity price, with exogenous shifts in electricity supply being attractive candidates. A similar approach is used by Fabra et al. (2021) in the Spanish electricity market.

Wind energy production made up almost 25% of the electricity mix in Germany in

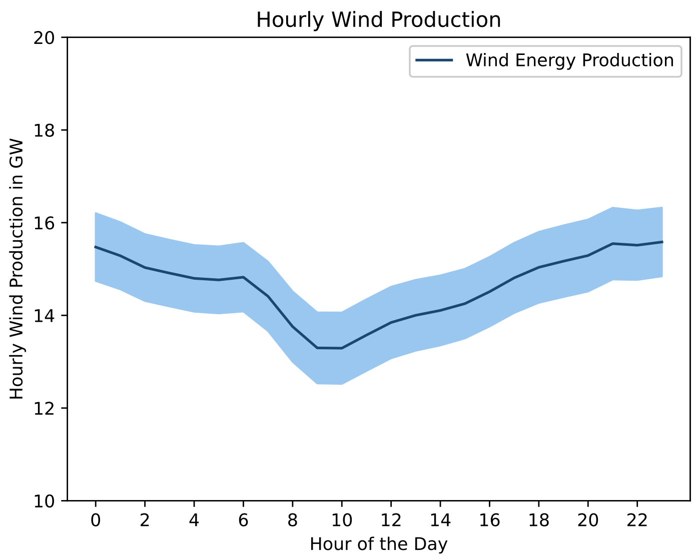

Figure 2.8. Wind energy Production Profile

This figure presents the hourly wind energy production profile over the day in GW between January 1, 2019 and December 31, 2020. The blue band denotes the 95% confidence interval.

2019.[15] Shifts in wind energy production are thus substantial drivers of overall electricity production and in turn the electricity price and are, conditional on the electricity price and weather conditions, unlikely to be correlated with domestic electricity consumption. Figure 2.8 shows the electricity production profile from wind energy. We observe slight variations in electricity production over the day. These variations are exclusively due to exogenous weather conditions. One possible explanation for variations in wind energy production would be the systematic switching on and off of wind turbines in response to the electricity price and grid loads. This would, of course, pose a problem for the exclusion restriction of our instrument. However, as part of the renewable-energies-law (*Erneuerbare-Energien-Gesetz EEG*), wind turbine operators are guaranteed to receive a fixed price for the electricity they feed into the grid. As the marginal costs of operating a wind turbine are very low, wind turbine operators have no incentive to ever actively shut down wind turbines, making variations in wind energy production a result of wind alone.

One could make an argument for weather conditions influencing wind energy production also influencing households electricity consumption, for example through summer and winter cycles, households' tendency to spend the day outside etc. We therefore also control for time fixed effects and local weather for each house-

15. Source: Fraunhofer-Institut, https://www.ise.fraunhofer.de/de/presse-und-medien/news/2019/oeffentliche-nettostromerzeugung-in-deutschland-2019.html, last visited: Jan 5, 2022

hold through daily temperature, sun hours and air pressure based on daily weather data provided by the weather station closest to the household. To estimate the sensitivity with respect to price changes, we propose a modification of the following model[16]:

$$\begin{aligned} p_{t,h} &= \alpha_h + \beta_{dow(t)} + \gamma_{m(t)} + \zeta T_t + \eta * S_t + \theta * P_t + \iota * W_h + \epsilon^1_{t,h} && (2.1)\\ y_{i,t,h} &= \kappa_h + \lambda_{dow(t)} + \mu_{m(t)} + \nu_i + \xi * T_{i,t} + \pi * S_{i,t} + \rho * P_{i,t} + \sigma * p_{t,h} + \epsilon^2_{i,t,h} && (2.2) \end{aligned}$$

$p_{t,h}$ denotes the electricity price at date t at hour h and is instrumented with wind energy production W in the first stage, as shown in equation (2.1). Wind energy production is measured in GW. The outcome variable $y_{i,t,h}$ denotes household i's electricity consumption in kWhs at date t in hour h.

We also include hour-of-the-day-, day-of-the-week-, samplemonth-, and household fixed effects. Finally, we control for local weather in household i's zipcode area using daily average temperature T in degree Celsius, hours-of-sunshine S per day, and average air pressure P (in hPA) on date t. The coefficient of interest is the coefficient σ in the second stage equation (2.2). It denotes the average marginal effect of increasing the EPEX day-ahead price by 1 ct per kWh, induced by fluctuations in wind energy production, on hourly electricity consumption.

Note that the first stage presented here does not exactly represent the first stage we will use in the final regressions. The reason for this is the panel structure of our data. The most intuitive 2SLS approach would simply estimate equation (2.1) on a time series of hourly EPEX prices between January 2019 and December 2020 and then plug in the predicted values for the EPEX prices into equation (2.2) (which is estimated on the whole panel), while adjusting the standard errors to reflect the correct residual variance estimator. However, we additionally include local weather, such as sun hours, air pressure and temperature on household-date-level in the second stage. Including them in the first stage is econometrically correct and necessary, but requires a separate first stage regression for each household panel. We thus estimate the following model on the whole panel using 2SLS:

16. Note that the specification laid out here does not exactly represent the specification we will ultimately use, but is rather meant to illustrate the intuition behind the 2SLS-approach used in the *actual* specifications (2.3) and (2.4)

$$p_{i,t,h} = \alpha_h + \beta_{dow(t)} + \gamma_{m(t)} + \delta_i + \zeta * T_{i,t} + \eta * S_{i,t} + \theta * P_{i,t} + \iota * W_{t,h} + \epsilon^1_{i,t,h} \tag{2.3}$$

$$y_{i,t,h} = \kappa_h + \lambda_{dow(t)} + \mu_{m(t)} + \nu_i + \xi * T_{i,t} + \pi * S_{i,t} + \rho * P_{i,t} + \sigma * p_{i,t,h} + \epsilon^2_{i,t,h} \tag{2.4}$$

That is, we treat the first stage as if it was determined at the household level. Hence, even though $p_{i,t,h} = p_{j,t,h}$ for all households i, j in our sample, we distinguish between households in the first stage and control for local weather phenomena as if the first stage was determined at the household level.

In a next step, we can exploit the within-day variation in the hourly electricity price to estimate intra-day price elasticities. In this approach, the identifying variation would only stem from the hour-to-hour variation in the electricity price, not from price variation across days. To do so, we include date fixed effects in regressions (2.3) and (2.4)[17]:

$$p_{i,t,h} = \alpha^w_h + \beta^w_t + \delta^w_i + \zeta^w * T_{i,t} + \eta^w * S_{i,t} + \theta^w * P_{i,t} + \iota^w * W_{t,h} + \epsilon^{1,w}_{i,t,h} \tag{2.5}$$

$$y_{i,t,h} = \kappa^w_h + \lambda^w_t + \nu^w_i + \xi^w * T_{i,t} + \pi^w * S_{i,t} + \rho^w * P_{i,t} + \sigma^w * p_{i,t,h} + \epsilon^{2,w}_{i,t,h} \tag{2.6}$$

Using this specification, we estimate the daily price reactions stemming from intra-day fluctuations in wind energy production. By adding date fixed effects, we exploit price variation that only occurs over the day, leaving us with intra-day identifying variation. This way, we can be confident to say that the price reactions we estimate are not long-run consumption shifts but actually short-term reactions to previously announced price variation induced by fluctuations in wind energy production. They are thus the statistics needed for effective demand side management.[18]

Finally, aggregating consumption on a daily level also allows to exploit the across-day variation in the electricity price. That is, we estimate equations (2.3) and (2.4), dropping the index h by aggregating electricity consumption on a daily level and using the average hourly electricity price over the day as the (endogenous) regressor. Aggregating consumption on a daily level thus eliminates all load shifting that happens within the day and instead focuses on load shifting across days. Equations (2.7) and (2.8) present the corresponding first- and second stage[19]:

17. We add the superscript w (for *within*) to indicate the difference to the coefficients in the model outlined before.

18. Note that short- and long-term price reactions here refer to intra- and inter-day price sensitivities and do not coincide with long- and short-run price sensitivities estimated in the common literature on price elasticities.

19. We add the superscript a (for *across*) to indicate the difference to the coefficients in the models outlined above.

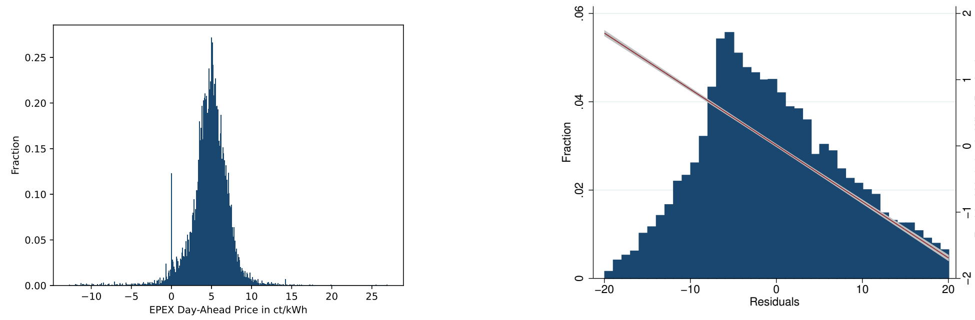

(a) Histogram of EPEX Prices

(b) Relevance of the First Stage

Figure 2.9. Price Variation and Effect of Wind on Price

Histogram of EPEX prices and first stage. Figure 2.9a shows that EPEX prices can be negative and bunch at 0. The histogram in 2.9b displays the density of residuals of wind energy production (on the left y-axis). These residuals stem from a regression of wind energy production on hour-, samplemonth- and day-of-the-week fixed effects. The red line plots a linear regression of the residual prices (based on a regression of the hourly electricity price on hour-, samplemonth- and day-of-the-week fixed effects) on the wind energy production residuals. The grey area denotes the 95% confidence interval. This graph is inspired by Dahl, Kostøl, and Mogstad (2014).

$$p_{i,t} = \beta^a_{dow(t)} + \gamma^a_{m(t)} + \delta^a_i + \zeta^a * T_{i,t} + \eta^a * S_{i,t} + \theta^a * P_{i,t} + \iota^a * W_{t,h} + \epsilon^{1,a}_{i,t} \quad (2.7)$$

$$y_{i,t} = \lambda^a_{dow(t)} + \mu^a_{m(t)} + \nu^a_i + \xi^a * T_{i,t} + \pi^a * S_{i,t} + \rho^a * P_{i,t} + \sigma^a * p_{i,t} + \epsilon^{2,a}_{i,t} \quad (2.8)$$

In Figure 2.9, we present evidence regarding the validity of our instrument. To show the relevance of the first stage, we first present a histogram of the EPEX day-ahead prices in panel (a). Prices are neither topcoded nor winsorized. Panel (b) presents the residual variation in EPEX day-ahead prices on the x-axis, relative to the residual variation in wind energy production, conditional on hour-, samplemonth- and day-of-the-week fixed effects, on the right y-axis. Residuals are cut off for values larger than 20 or smaller than -20. We first see that both without and with conditioning on time- and individual fixed effects there is considerable variation in electricity prices. Secondly, we clearly see the highly significant, negative relationship between wind energy production and the electricity price, illustrated by the solid red line in panel (b).

2.5.3 Estimating Individual Price Reactions

The panel structure of our data allows us to additionally estimate individual household price reactions. To see this, we can rewrite equations (2.3) and (2.4) as follows for each household i[20]:

20. We add the superscript i (for *individual*) to indicate the difference to the coefficients in the models outlined above.

$$p_{t,h} = \alpha_h^i + \beta_{dow(t)}^i + \gamma_{m(t)}^i + \zeta^i * T_t + \eta^i * S_t + \theta^i * P_t + \iota^i * W_{t,h} + \epsilon_{t,h}^{1,i} \quad (2.9)$$
$$y_{t,h} = \kappa_h^i + \lambda_{dow(t)}^i + \mu_{m(t)}^i + \xi^i * T_t + \pi^i * S_t + \rho^i * P_t + \sigma^i * p_{t,h} + \epsilon_{t,h}^{2,i} \quad (2.10)$$

Where the notation stays the same as before, but the index i and the corresponding fixed effect is dropped. Intuitively, we now estimate σ separately on 899 different time series instead of a panel containing 899 households.

2.6 Results

In order to get a better grasp of our results, we first present results from a naive OLS regression of electricity consumption on electricity prices as a baseline. Of course, these results are expected to be biased due to the reversed causality and omitted variable bias discussed earlier. In particular, consumption patterns shown by most households affect both electricity consumption of the households in our sample and the electricity price set at the European Power Exchange. OLS should thus underestimate the magnitude of price reactions, especially if we do not control for hourly and seasonal consumption patterns that affect the electricity price the most.

2.6.1 OLS Estimates

Table 2.2 reports results from simple OLS estimations of hourly electricity consumption on the hourly electricity price per kWh. In each specification, we subsequently add fixed effects and local weather controls. Finally, in the last specification, we run the regression on a different sample of households, namely, the control households from a study conducted between January 2019 and February 2020 in Germany. This dataset includes the consumption data of 130 households in 15-minute intervals as well as state-level geolocation data. Using the center point of the respective state and merging the weather data of the closest weather station of that center point, we can run the same set of regressions on this separate dataset.

In our main sample, without controlling for hour- or month fixed effects, estimation results report a significantly positive coefficient for the EPEX spot price on hourly electricity consumption. Without controlling for time fixed effects, an increase in the hourly price of 1 ct/kWh is associated with an average increase in hourly electricity consumption of .006 to .012 kWh. After controlling for hour-of-the-day fixed effects and local weather, this association becomes significantly negative. Together, the results imply a strong positive bias of simple OLS estimates, that is partially corrected when controlling for time and local weather, which are closely related to both the electricity price and wind production.

Running the same OLS regression on the set of control households returns a positive, significant estimate for the electricity price effect, while the coefficients on local weather effects work in the same direction and stay significant. Note, however, that their absolute magnitude is lower as average consumption in the control sample is significantly lower. Table 2.C.1 in the appendix presents the whole OLS regression table for the control sample. The change in coefficients after adding fixed effects hints at the influence of daily- and seasonal household consumption patterns.

2.6.2 IV Estimates

We now turn to the causal analysis of electricity price sensitivities. Table 2.B.1 in the appendix delivers summary statistics of the electricity price, wind production and weather controls. Table 2.3 presents the IV estimation results from regression equations (2.3) and (2.4). In column 1, no fixed effects are included, in column 2, we add day-of-the-week fixes effects, column 3 adds hour-of-the-day fixed effects, column 4 adds sample month fixed-effects and in column 5, we additionally control for local weather. We see that estimates are significantly negative across all specifications, confirming our initial suspicion that OLS estimates may be biased due to the reversed causality between consumption patterns and electricity prices. The change in the reported estimates between columns 4 and 5 underlines the importance of time trends in our sample. With variations in wind production and aggregate electricity demand over time, we see that ignoring these time trends would produce biased results. After controlling for sample month fixed-effects, estimates are consistent across specifications and significantly negative. In our preferred specification in column 6, we find that an increase in the electricity price of 1 ct/kWh decreased hourly electricity consumption by .016 kWh, or approximately 3% of hourly electricity consumption on average. This estimate is highly significant. With an average (after taxes and levies) electricity price of 23.9 ct/kWh (in our sample) and an average hourly consumption of .530 kWhs, this translates into a price elasticity of $-0.016 \times 23.9/0.53 \approx -0.676$.
In contrast, running the same analysis on the sample of the control households returns a statistically insignificant, positive coefficient on the EPEX price. The effect of local weather patterns points in the same direction as before, but shows a smaller magnitude, which is partially explained by the lower consumption baseline and different geolocation data level in the control sample.

To illustrate our results and to check for potential nonlinearlities in the relationship between the electricity price and electricity consumption, we present a residual plot in Figure 2.10. We plot the residuals of a regression of hourly electricity consumption on household-, hour-of-the-day-, samplemonth-, and day-of-the-week fixed effects as well as local weather controls against the predicted price, based on a regression of the EPEX spotprice on wind production, hour-of-the-day, samplemonth- and day-of-the-week fixed effects as well as local weather controls. We then fit an OLS model of residualized consumption on the predicted and squared predicted price. We can clearly see the negative relationship between the predicted price and consumption residuals. Including the squared predicted price reveals that households react stronger to changes in the electricity price if the price is low.

Table 2.2. Effect on Electricity Consumption: OLS Estimates

Dependent Variable	**Hourly Electricity Consumption** (1)	(2)	(3)	(4)	(5)	(6)	(Control)
EPEX Spot Price	.007***	.007***	.012***	.006***	-.013***	-.015***	.002***
	(.002)	(.002)	(.002)	(.002)	(.002)	(.002)	(.001)
Daily Temperature						-.081***	-.012*
						(.007)	(.006)
Hours of Sunshine						-.050***	-.005*
						(.004)	(.003)
Air Pressure						-.005***	-.004**
						(.001)	(.002)
Household fixed-effects	No	Yes	Yes	Yes	Yes	Yes	Yes
DOW FEs	No	No	Yes	Yes	Yes	Yes	Yes
Hour FEs	No	No	No	Yes	Yes	Yes	Yes
Sample Month fixed-effects	No	No	No	No	Yes	Yes	Yes
Local Weather Controls	No	No	No	No	No	Yes	Yes
R^2_{within}	0.0003	0.0003	0.0017	0.0075	0.0370	0.0388	0.0825
$R^2_{between}$	0.0474	0.0474	0.0458	0.0461	0.0548	0.0606	0.0192
$R^2_{overall}$	0.0005	0.0005	0.0019	0.0069	0.0355	0.0367	0.0703
Observations	5’583’500	5’583’500	5’583’500	5’583’500	5’583’500	5’360’395	1’102’116
Households	829	829	829	829	829	796	116

Notes: For readability, local weather controls (measured in degrees celsius, hours, and hPA, respectively) are divided by 10. That is, the effect on hourly electricity consumption of an increase in temperature by one degree celsius is $\hat{\beta}_{temp}/10$

Results stem from OLS regression of hourly electricity consumption on the hourly electricity price per kWh, based on equation (2.4), without instrumenting for the hourly electricity price. We subsequently add household-, day-of-the-week-, hour-of-the-day- and sample month fixed effects. Finally, we add local weather controls in column (6), based on the weather data provided by the closest weather station. In the last column, we run the same regression displayed in column (6) on the control sample. Standard errors are clustered at the household level.

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

Table 2.3. Effect on Electricity Consumption: IV Estimates

Dependent Variable	**Hourly Electricity Consumption**						
	(1)	(2)	(3)	(4)	(5)	(6)	(Control)
EPEX Spot Price	-.073***	-.072***	-.074***	-.071***	-.014***	-.016***	.001
	(.000)	(.004)	(.004)	(.004)	(.002)	(.001)	(.001)
Daily Temperature						-.081***	-.013**
						(.007)	(.006)
Hours of Sunshine						-.050***	-.004*
						(.004)	(.002)
Air Pressure						-.005***	-.003**
						(.001)	(.002)
Household fixed-effects	No	Yes	Yes	Yes	Yes	Yes	Yes
DOW FEs	No	No	Yes	Yes	Yes	Yes	Yes
Hour FEs	No	No	No	Yes	Yes	Yes	Yes
Sample Month fixed-effects	No	No	No	No	Yes	Yes	Yes
Local Weather Controls	No	No	No	No	No	Yes	Yes
R^2_{within}	-	-	-	-	0.0370	0.0388	0.0824
$R^2_{between}$	-	0.0474	0.0472	0.0469	0.0548	0.0606	0.0194
$R^2_{overall}$	-	0.0005	0.0008	0.0005	0.0355	0.0367	0.0703
Observations	5’583’500	5’583’500	5’583’500	5’583’500	5’583’500	5’360’395	1’102’116
Households	829	829	829	829	829	796	116

Notes: For readability, local weather controls (measured in degrees celsius, hours, and hPA, respectively) are divided by 10. That is, the effect on hourly electricity consumption of an increase in temperature by one degree celsius is $\hat{\beta}_{temp}/10$
Results stem from an IV regression of hourly electricity consumption on the hourly electricity price per kWh, based on equation (2.4), with the hourly electricity price being instrumented for using hourly wind production in Germany according to equation (2.3). We subsequently add household-, day-of-the-week-, hour-of-the-day- and sample month fixed effects. Finally, we add local weather controls in column (6), based on the weather data provided by the closest weather station. In the last column, we run the same regression displayed in column (6) on the control sample. Standard errors are clustered at the household level.
*** $p < 0.01$ ** $p < 0.05$ * $p < 0.1$

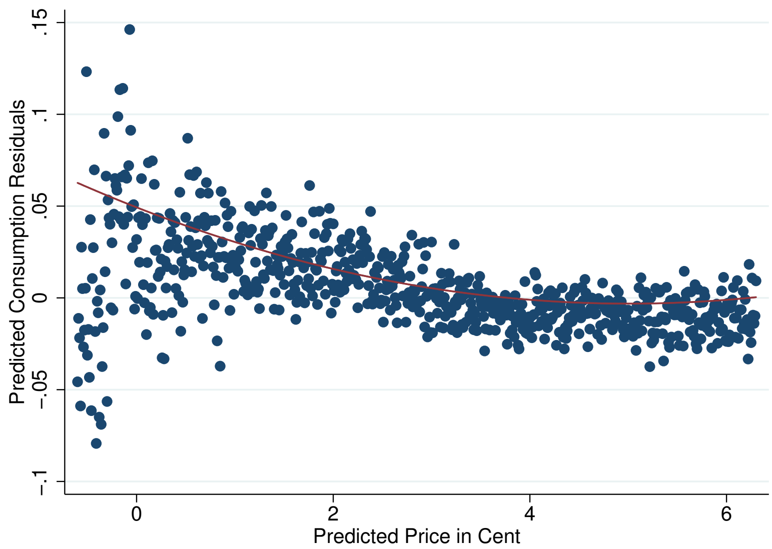

Figure 2.10. Scatterplot of Predicted Price and Predicted Consumption Residuals

This figure presents a scatterplot of residualized consumption from a regression of consumption on local weather controls, hour-, samplemonth-, day-of-the-week- and household-fixed effects against predicted price bins, based on a regression of the electricity price on the same set of covariates plus wind energy production. The red line plots fitted consumption residuals, based on a regression of residualized consumption on the predicted and squared predicted price. Note that each dot represents an interval of .01 cent on the x-axis, with the y-values being collapsed within that interval to keep the number of dots manageable.

A similar pattern can be found if we analyze the reaction to price changes for different hours of the day. We run the same IV estimation outlined in equations (2.3) and (2.4) separately for each hour of the day.[21] Figure 2.11, panel (a) displays the price effect profile over the day. We observe the largest price reactions during the night and afternoon, when overall consumption is low. This may seem counterintuitive given that the largest scope for consumption adjustment can be found when consumption is high, but it is in line with our findings in Figure 2.10: Households are more price sensitive if the baseline price is low. For high prices, which are prevalent during the morning- and evening hours (see Figure 2.4), consumers react less to changes in the price, indicating that the relative size of the price change is an important determinant of the price sensitivity. Figure 2.C.1 in the appendix reveals that we do not find a similar pattern in the control sample - the estimated price sensitivity remains small and insignificant over the day.

Figure 2.11, panel (b) reveals that the price effect pattern over the day is driven by the reduced form, which varies by a factor larger than 4 over the day, instead of the first stage, which only varies between −.08 and −.12. To see this relationship, recall that the classic 2SLS estimator can be written as the ratio of the effect of the

21. Naturally, we leave out hour-of-the-day fixed effects in this specification.

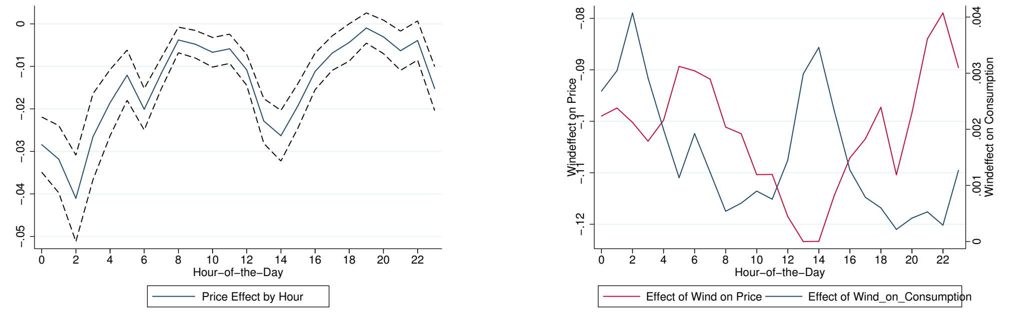

(a) Price Effect by Hour-of-the-Day

(b) Association of Wind Production and Electricity Consumption

Figure 2.11. Price Effect and Association of Wind Production and Electricity Consumption

Figure 2.11a presents the estimated price effect for each hour of the day. The dashed line denotes the 95%-confidence interval. Figure 2.11b shows that this effect is driven by the reduced form, not the first stage. The red line represents the first stage results over the day, measured on the left y-axis. The blue line represents the reduced form results, measured on the right y-axis. The specifications include local weather controls and individual-, sample month- and day-of-the-week-fixed effects. Standard errors are clustered at the household level.

instrument on the outcome (the reduced form) to the effect of the instrument on the endogenous regressor (the first stage).

$$\begin{aligned} \hat{\beta}_{IV} &= \left(z'x\right)^{-1} z'y \\ &= \frac{\left(z'z\right)^{-1} z'y}{\left(z'z\right)^{-1} z'x} \\ &= \frac{dy/dz}{dx/dz} \end{aligned} \tag{2.11}$$

Two explanations may arise as to why price reactions are strongest during the night and afternoon. Intuitively, we would expect price reactions to be largest when consumption is highest, that is, during the morning and evening hours, because that is when the scope for load shifting is largest. However, households may simply see no scope for consumption reduction during the high price hours, as workers go to work in the morning and come back in the evening, with little wiggle room for time adjustments around those particular hours, whereas nighttime and afternoon consumption can be easier shifted across hours. Another explanation relates to Figure 2.10: Households may be more attentive to price changes, the larger the price change is compared to the electricity price. Thus, the price sensitivity would be driven by the baseline electricity price instead of baseline consumption. To find out which factor drives our results, we run the same analysis separately for each day of the week. The results are presented in Figure 2.12. We find households to be more price sensitive during the weekend than on regular days. That is, we find stronger

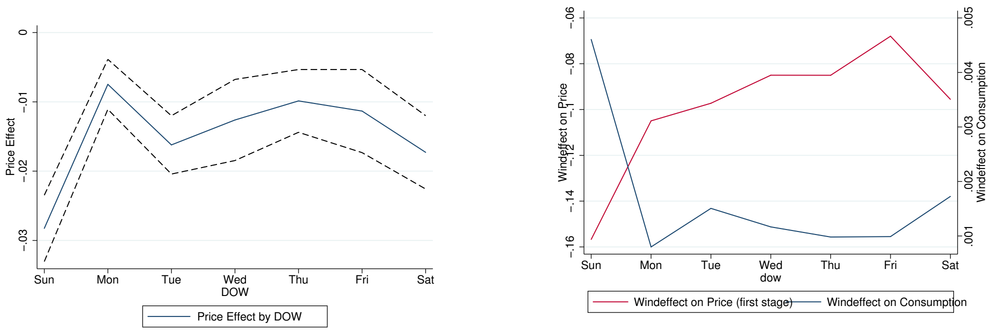

(a) TE by Day-of-the-Week

(b) Association of Wind Production and Electricity Consumption

Figure 2.12. Price Effect and Association of Wind Production and Electricity Consumption

Figure 2.12a presents the estimated price effect over the week. The dashed line denotes the 95%-confidence interval. Figure 2.12b shows that again, the effect is driven by the reduced form instead of the first stage, with its relative magnitude being significantly larger. As before, the red line represents the first stage results over the week, measured on the left y-axis. The blue line represents the reduced form results, measured on the right y-axis. The specifications include local weather controls and individual-, sample month- and hour-of-the-day-fixed effects. Standard errors are clustered at the household level.

effects when baseline consumption is high and prices are low, indicating that indeed the low baseline price is decisive for the price sensitivity and baseline consumption may not the driving factor of its magnitude. Figure 2.C.2 in the appendix reveals that we do not find a similar pattern in the control sample.

2.6.2.1 Price Reactions for Different Time Horizons

So far, we cannot inform whether electricity price changes induce overall consumption patterns to change across days and weeks or whether the price effect in electricity consumption is driven by hourly adjustment of consumption. That is, do households rather react to expensive days by shifting electricity consumption from one day to another, or do they react to expensive hours of the day by shifting electricity consumption to cheaper hours of the day?

Intra-Day Price Effects. To answer this question, we zoom in on our analysis by including even finer grained time fixed effects. Instead of sample month fixed effects, we now use date fixed effects as illustrated in equations (2.5) and (2.6) and thus assess only the effect of within-day price variation. Table 2.4 reports the estimated price sensitivities. While the estimated price sensitivity in the control households remains unchanged, the estimated intra-day price effect in our sample is more than twice the size of the overall estimated price effect in Table 2.3. In our preferred specification in column (6), we report that an increase in the hourly electricity price of one Cent significantly decreases hourly electricity consumption by 0.037 kWhs, or 7% on average. Our estimates imply an intra-day price elasticity of $-0.037 \times 23.9/0.53 = -1.668$. This is especially interesting, because it suggests that households react even stronger to short-term electricity price fluctuations over the day than fluctuations across days.

Inter-Day Price Effects. In order to assess the possibility of electricity load shifting over several days, we can also estimate price effects on daily level data. That is, we aggregate daily consumption and estimate the model outlined in equations (2.7) and (2.8). By aggregating consumption on the daily level, this approach ignores any within-day price variation and thus any intra-day load shifting and instead focuses purely on the price sensitivity linked to electricity load shifting across days, for example through appliances whose usage can be adjusted across days (such as washing machines or dryers). Table 2.5 presents the results. In our preferred specification in column (5), we estimate that an increase in the average hourly electricity price over the day of one Cent significantly decreases daily electricity consumption by .241 kWhs, or 1.8% on average. Using the average daily electricity consumption of 12.71 kWhs, our estimate for the price effect implies an inter-day price elasticity of $-.241 \times 23.9/12.710 = -0.453$, which is significantly smaller than the overall price elasticity based on the original price sensitivity estimated on the hourly dataset and the calculated intra-day price elasticity.

Two possible explanations arise: First, the scope for intra-day price reactions is larger than inter-day price reactions. Some activities, such as cooking or showering, can, to a certain extend, be shifted within a day. A person can decide to shower in the morning instead of the evening, or to cook at 5 p.m. instead of 7 p.m., but the scope for load shifting across several days is limited for many such activities. A second explanation comes to mind when we take a closer look at the information channels for customers. As we showed in Section 2.3, customers can see the electricity prices for the next day at 2 p.m. on the electricity provider's website. This means that from 2 p.m. it is possible to plan electricity consumption for the next day, whereas consumption planning further in the future is much harder, as households cannot reliably plan with electricity prices across days.[22] Price security within a day allows households to easily shift loads over the day, but meaningful inter-day load shifting with knowledge of future prices is only possible after 2 p.m.

The policy implications of this result are large: The general consensus of the previous literature on electricity price elasticities, both in households and businesses was that short-term price elasticities are much lower than long-run elasticities, to the point of being insignificant for policymakers. This often led to the conclusion that intra-day price variation is not an effective tool for demand side management. However, it is important to note that this conclusion is confusing two different concepts, namely short-term- and intra-day price elasticities. The former describes intra- and inter-day price elasticities in the short-term window after TVP introduction, the lat-

22. Though they can as least use the general price pattern described in Figure 2.4.

Table 2.4. Effect on Electricity Consumption: IV Estimates of Intra-Day Price Effects

Dependent Variable	**Hourly Electricity Consumption**						
	(1)	(2)	(3)	(4)	(5)	(6)	(Control)
EPEX Spot Price	-.073***	-.072***	-.069***	-.015***	-.037***	-.037***	.001
	(.000)	(.004)	(.003)	(.002)	(.002)	(.003)	(.001)
Daily Temperature						-.101***	.005
						(.019)	(.018)
Hours of Sunshine						-.055***	-.003
						(.005)	(.004)
Air Pressure						-.005	-.005
						(.008)	(.010)
Household fixed-effects	No	Yes	Yes	Yes	Yes	Yes	Yes
Hour FEs	No	No	Yes	Yes	Yes	Yes	Yes
Sample Month fixed-effects	No	No	No	Yes	Yes	Yes	Yes
Date fixed-effects	No	No	No	No	Yes	Yes	Yes
Local Weather Controls	No	No	No	No	No	Yes	Yes
$R^2_{centered}$	-	-0.0362	-0.0249	0.0011	0.0006	0.0012	0.0001
$R^2_{uncentered}$	-	-0.0362	-0.0249	0.0011	0.0006	0.0012	0.0001
Observations	5’583’500	5’583’500	5’583’500	5’583’500	5’583’500	5’360’395	1’102’116
Households	829	829	829	829	829	796	116

Notes: For readability, local weather controls (measured in degrees celsius, hours, and hPA, respectively) are divided by 10. That is, the effect on hourly electricity consumption of an increase in temperature by one degree celsius is $\hat{\beta}_{temp}/10$

Results stem from IV estimation of hourly electricity consumption on the hourly electricity price per kWh, based on equation (2.6), with the hourly electricity price being instrumented for using hourly wind production in Germany according to equation (2.5). We subsequently add household-, hour-of-the-day-, sample month- and date-fixed effects. Finally, we add local weather controls in column (6), based on the weather data provided by the closest weather station. In the last column, we run the same regression displayed in column (6) on the control sample. Standard errors are clustered at the household level.

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

ter describes the price elasticity following sudden price variation over the day. Our estimation results help us to separate both concepts, as we estimate the latter.

Overall, we find highly significant demand responses to electricity price changes for each hour of the day and each day of the week, with the magnitude of the estimated price sensitivity changing significantly over the day and week. Our results indicate that households react stronger to price changes if the baseline price is low. Additionally, we find that the resulting price elasticity with respect to hourly changing electricity prices over the day is significantly larger than the price elasticity with respect to daily changing electricity prices.

These results are robust with respect to various specifications and sampling: In Appendix 2.D, we report the estimated intra- and inter-day price sensitivities for a series of robustness checks. We find that our results are robust to the exclusion of observations during the night, confirming that the estimated negative price sensitivity is not mainly driven by excessive nighttime consumption. In order to assess the role of staying at home for the price elasticity of demand, we also run the analyses separately on observations before the first lockdowns of the COVID-19 pandemic and during the pandemic. With the higher average socioeconomic status of the zipcodes in our sample, we expect that household members in our sample are more likely to work in home office compatible jobs with an increased possibility of staying at home, which potentially increases the scope for load shifting over the day - especially during weekdays. While we do find slight changes in the estimated price sensitivity before and during the pandemic, these differences are extremely small. We also estimate households to be significantly more price sensitive on weekends than on weekdays. While the estimated price sensitivity in the control sample is small and insignificant for both weekends and weekdays, we find that in our sample, both the intra- and the inter-day price elasticities on weekends are almost twice as large as the price elasticities during the week. Including all clusters of households in our analysis increases the magnitude of the estimated price sensitivity as well as the resulting price elasticities. Including only the "low consumption"-cluster decreases the magnitude of the estimated price sensitivity, but - due to the lower baseline consumption - results in a larger estimated price elasticity. We also show that our results are robust to using the wind energy production day-ahead forecast as the instrument (instead of *actual* wind energy production). Finally, in Appendix 2.E, we use the difference in hourly electricity consumption, the hourly electricity price and hourly wind production and their respective 24-hour lags as the outcome, regressor, and instrument, respectively in equations (2.3) and (2.4). We thus address the endogeneity inherent to the EPEX prices and include daily and seasonal time trends indirectly. Using this approach, we estimate a slightly smaller (in absolute terms), yet still significantly negative price sensitivity for the households in our sample, and a small, positive price sensitivity for the households in the control sample.

Table 2.5. Effect on Daily Electricity Consumption: IV Estimates

Dependent Variable	**Daily Electricity Consumption**					
	(1)	(2)	(3)	(4)	(5)	(Control)
EPEX Spot Price	-1.838***	-1.837***	-1.903***	-.187***	-.241***	.026
	(.093)	(.093)	(.100)	(.037)	(.035)	(.026)
Daily Temperature					-1.855***	-.313**
					(.170)	(.136)
Hours of Sunshine					-1.195***	-.100*
					(.099)	(.058)
Air Pressure					-.133***	-.086**
					(.033)	(.042)
Household fixed-effects	No	Yes	Yes	Yes	Yes	Yes
DOW FEs	No	No	Yes	Yes	Yes	Yes
Sample Month fixed-effects	No	No	No	Yes	Yes	Yes
Local Weather Controls	No	No	No	No	Yes	Yes
R^2_{within}	0.0021	-	-	0.1449	0.1523	0.0709
$R^2_{between}$	0.0462	0.0462	0.0459	0.0529	0.0581	0.0193
$R^2_{overall}$	0.0027	0.0027	0.0047	0.0993	0.1026	0.0362
Observations	232’992	232’992	232’992	232’992	223’683	45’936
Households	829	829	829	829	796	116

Notes: For readability, local weather controls (measured in degrees celsius, hours, and hPA, respectively) are divided by 10. That is, the effect on daily electricity consumption of an increase in temperature by one degree celsius is $\hat{\beta}_{temp}/10$
Results stem from IV estimation of daily electricity consumption on the average hourly electricity price per kWh over the day, based on equation (2.8), with the average hourly electricity price over the day being instrumented for using average hourly wind production over the day in Germany according to equation (2.7). We subsequently add household-, day-of-the-week- and sample month-fixed effects. Finally, we add local weather controls in column (5), based on the weather data provided by the closest weather station. In the last column, we run the same regression displayed in column (5) on the control sample. Standard errors are clustered at the household level.
*** $p < 0.01$ ** $p < 0.05$ * $p < 0.1$

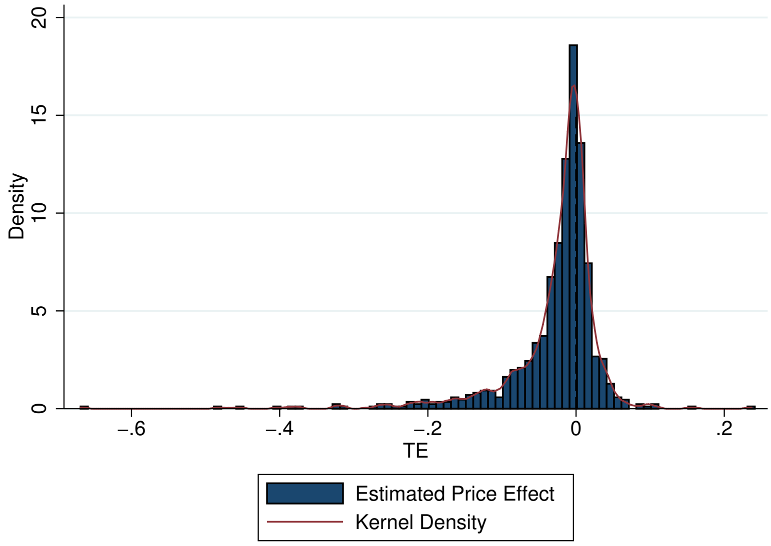

Figure 2.13. Distribution of Individual Price Effects

This figure presents a histogram and kernel density of estimated household level price sensitivities estimated according to equations (2.9) and (2.10). That is, for this figure, we estimate a separate overall price sensitivity for each household in our sample (including all clusters of households).

2.6.3 Individual Price Sensitivities

So far, we have estimated the electricity price effect on consumption on the whole sample. However, estimating one parameter for the whole population ignores the potential heterogeneity in the parameter of interest, which can be defined at the individual level. As we can see in Figure 2.6, electricity consumption is subject to substantial heterogeneity, with not only the absolute magnitude in consumption, but also consumption patterns differing between households.

The panel structure of our dataset allows us to estimate individual-level price sensitivities for each of the 899 households in our sample. To do so, we estimate the model outlined in equations (2.9) and (2.10). For this approach, we include all 899 households in our sample and estimate 899 separate electricity price sensitivities. Figure 2.13 presents the distribution of estimated price sensitivities as well as kernel density plots. We observe the sensitivities to be centered close to zero, though with significantly more probability mass in the negative range. We can thus say that the negative price sensitivity estimated earlier is not driven by a few highly price elastic households or households using an IM, but by a large portion of the households in our sample.

2.7 Conclusion

Residential households' sensitivity to the electricity price is crucial information to evaluate the effectiveness of time-varying electricity prices. This paper provides empirical estimates of households' real-time price elasticities of electricity demand. We use real-world high-frequency data on households' electricity consumption along with hourly electricity prices, provided by an electricity provider operating in Germany and Austria, which passes hourly electricity prices determined at the European Power Exchange *EPEX* on to its customers, thus providing households with a real-time-pricing tariff. Applying a simple K-means clustering algorithm to the sample of the provider's customers reveals a cluster of households apparently using an automated demand response to low prices during the night.

Using hourly wind production as an instrument for the electricity price in a 2SLS framework to overcome the endogeneity inherent to simple OLS due to the EPEX pricing mechanism, we estimate a significantly negative overall household price elasticity of −0.69. Our results additionally reveal important nonlinearities in the housholds' demand response to price variations. While demand is fairly price elastic for low electricity prices, we find that demand responses to price changes are smaller in magnitude when prices are high. These benchmark results show that households in our sample are indeed highly price elastic and that, if real-time pricing is adapted, hourly changing electricity prices can be an successful strategy to control electricity consumption.

By exploiting the within-day price variation in the electricity price, we find an even larger intra-day price elasticity of −1.668. In contrast, aggregating our data on a daily level, we find a significantly smaller inter-day price elasticity of −0.453. That is, households in our sample react more strongly to price variation over the day than price variation across days by shifting electricity consumption within the day rather than across days. These important results indicate that load shifting of activities that usually have to be done each day (e.g. cooking, showering etc.) is an important driver of the households' demand response to price changes. This has important welfare implications for policymakers: While dynamic electricity price tariffs are often discarded as unreasonable and inexpedient (Forsa (2015)), we show that in our sample, households not only strongly react to time-varying prices, but they do so even more stongly in a short time horizon. Conversely, while we still observe a highly significant demand response to price changes across days, this response is significantly smaller. Policymakers should thus consider the time frame if they hope real-time pricing to achieve electricity load shifting: continuously high prices over several days have limited potential of shifting electricity over time and should thus be handled with care, while short-term grid-wide consumption spikes can be partially absorbed by load shifting induced by real-time pricing.

Moreover, as shown in Appendix 2.F, we can show that, for mean-preserving hourly changing electricity prices, risk-neutral households are weakly better off than for constant electricity prices: For a given mean price, flexible electricity prices increase the consumer surplus of risk-neutral households. Due to the mean-preserving spread of electricity prices and risk-neutrality, the consumer surplus from fixed prices is the same as the surplus from flexible prices if households were forced to consume the same amount of energy they would have used if prices were fixed at the mean price. But allowing households to react and optimize their electricity consumption with respect to the electricity price makes households better off compared to the fixed price case. Similar mechanisms of welfare-increasing dynamic pricing have been shown, for example, in the context of airfares (Williams (2021)).

Finally, robustness checks excluding the nighttime hours between 10 p.m. and 6 a.m. reveal that our estimates are not driven by excessive nighttime consumption. We also find that households are more price elastic during weekends than during weekdays, which we believe is due to household members being at home and having a larger scope and more possibilities to adjust consumption. Interestingly, however, we do not find that households became more price elastic over the course of the COVID-19 pandemic and the consequent widespread roll-out of working from home. We also find that our results are robust with respect to different specifications (such as using the difference in 24-hour lags of the outcome, regressors and instrument as the outcome, regressors and instrument), sample compositions (including all household clusters or only the "low consumption"-cluster) and instruments (using the wind production forecast instead of actual wind production as the instrument).

As a cautionary note, it is important to acknowledge that the households in our sample most likely do not represent the average German household. Our results can thus not easily be generalized to a larger population. Using zipcode level information on sociodemographics, we find that the zipcodes in our sample exhibit a higher average socioeconomic status, a higher share of homeowners and families, a lower share of non-Germans as well as household heads over the age of 60 or under 30. However, except for the variables determining the socioeconomic status and the share of homeowners and families, these differences are small in absolute terms. Thus, while we cannot say much about the sociodemographics of the households in our sample, we can present evidence regarding their sociodemographic and socioeconomic environment. Note, however, that for a meaningful interpretation of our results on the policy implications of real-time pricing, households in our sample do not need to be representative. With the planned roll-out of smart meters in Europe, offering RTP to households and letting them self-select into treatment becomes a viable option - and a large share of potential customers are open to the idea of time-varying prices (Forsa (2015)).

This paper contributes to the still growing literature on real-time electricity demand price elasticities. The scarcity of such tariffs results in few data sources and only a handful of papers explicitly addressing dymanic electricity price elasticities. Previous work on this front is mostly based on experimental data (Wolak (2011), Allcott (2011), the only exception we are aware of being Fabra et al. (2021)) and thus seldom analyzes the price sensitivities to actually implementable pricing schemes. Our empirical analysis hopefully provides a base for future research to assess the effectiveness of dynamic electricity pricing schemes in different contexts and time periods and to further evaluate on the welfare implications of real-time pricing.

Appendix 2.A Representativeness of the Sample

To investigate the external validity of our results, we analyze the representativeness of our sample with respect to several sociodemographics and socioeconomic variables. Unfortunately, we do not have household-level information on these variables of interest for the households in our sample. However, we use information on sociodemographic variables on the zipcode level, provided by the data provider *microm*. Thus, while we do not have household-level information we have information about the socioeconomic environment each household lives in, which allows us to analyze whether a household's environment, which is a proxy for household characteristics, is representative of all households in Germany. In this section, we provide information on sociodemographic averages on zipcode level as well as distributions on income, family structures and socioeconomic status.

Table 2.A.1. Representativeness with Respect to Sociodemographics

Characteristic	Germany	Sample	p-value
Share of Homeowners	46.589	57.067	0.000
Share of HHs with A-level	23.551	25.042	0.000
Share of HHs with electric or hybrid cars	0.716	0.720	0.662
Share of HHs with non-German heads	12.287	10.39	0.000
Share of HHs with head under 30 years old	20.172	18.912	0.000
Share of HHs with head over 60 years old	34.061	33.609	0.003

Notes: This table shows the average share (in percent) of households falling into certain characteristics categories in Germany and in the zipcode areas the households in our sample live in. The displayed averages are generated as weighted averages, with the weights for Germany being calculated as the zipcode population divided by the total population of households in Germany ($\approx$ 41 million). For the sample, the weights are calculated as the number of households living in a zipcode area divided by the total number of households in our sample. The p-values stem from two-sided t-tests using the described weights.

Table 2.A.2. Representativeness with Respect to Income

Characteristic	Germany	Sample	p-value
Share of HHs w/ Income less than 1’100EUR p. Month	0.105	0.075	0.000
Share of HHs w/ Income between 1’100 and 1’500EUR p. Month	0.099	0.076	0.000
Share of HHs w/ Income between 1’500 and 2’000EUR p. Month	0.129	0.110	0.000
Share of HHs w/ Income between 2’000 and 2’600EUR p. Month	0.139	0.132	0.000
Share of HHs w/ Income between 2’600 and 4’000EUR p. Month	0.231	0.254	0.000
Share of HHs w/ Income between 4’000 and 7’500EUR p. Month	0.237	0.283	0.000
Share of HHs w/ Income higher than 7’500EUR	0.056	0.067	0.000

Notes: This table shows the average share (in percent) of households falling into different income categories for Germany and for our sample. The averages are generated as weighted averages, with the weights for Germany being calculated as the zipcode population divided by the total population of households in Germany ($\approx$ 41 million). For the sample, the weights are calculated as the number of households living in a zipcode area divided by the total number of households in our sample. The p-values stem from two-sided t-tests using the weights described earlier.

Table 2.A.3. Representativeness with Respect to Family Structures

Characteristic	Germany	Sample	p-value
Share of HHs falling into family structure category 1	11.109	5.130	0.000
Share of HHs falling into family structure category 2	11.109	6.774	0.000
Share of HHs falling into family structure category 3	11.109	8.185	0.000
Share of HHs falling into family structure category 4	11.109	9.729	0.000
Share of HHs falling into family structure category 5	11.109	11.381	0.216
Share of HHs falling into family structure category 6	11.109	12.851	0.000
Share of HHs falling into family structure category 7	11.109	13.973	0.000
Share of HHs falling into family structure category 8	11.109	15.049	0.000
Share of HHs falling into family structure category 9	11.119	16.920	0.000

Notes: This table shows the average share (in percent) of households falling into different family structure categories for Germany and for our sample. The averages are generated as weighted averages, with the weights for Germany being calculated as the zipcode population divided by the total population of households in Germany ($\approx$ 41 million). For the sample, the weights are calculated as the number of households living in a zipcode area divided by the total number of households in our sample. The p-values stem from two-sided t-tests using the weights described earlier. The categories 1-9 are defined as percentiles in the German population of households, which explains why the shares falling into each category are the same. This categorization results in the following thresholds for the share of single-households: Category 1: 16.95%, Category 2: 15.05%, Category 3: 13.86%, Category 4: 12.85%, Category 5: 11.63%, Category 6: 9.11%, Category 7: 9.02%, Category 8: 6.60%, Category 9: 4.90%

Table 2.A.4. Representativeness with Respect to Children

Characteristic	Germany	Sample	p-value
Share of HHS in Children Category 1	5.137	3.762	0.000
Share of HHS in Children Category 2	11.867	12.927	0.000
Share of HHS in Children Category 3	11.850	11.703	0.289
Share of HHS in Children Category 4	11.866	10.533	0.000
Share of HHS in Children Category 5	11.850	11.147	0.000
Share of HHS in Children Category 6	11.863	11.687	0.010
Share of HHS in Children Category 7	11.855	12.129	0.000
Share of HHS in Children Category 8	11.853	13.032	0.000
Share of HHS in Children Category 9	11.857	13.079	0.000

Notes: This table shows the average share (in percent) of households falling into different "share of children per household"-categories for Germany and for our sample. The averages are generated as weighted averages, with the weights for Germany being calculated as the zipcode population divided by the total population of households in Germany ($\approx$ 41 million). For the sample, the weights are calculated as the number of households living in a zipcode area divided by the total number of households in our sample. The p-values stem from two-sided t-tests using the weights described earlier. The categories 1-9 are defined by the average share of children in the households in a zipcode area. The thresholds are kept confidential by microm, but the categories are defined such that, with the exception of category 1, the share of households in each category is roughly the same for Germany. The categories are thus defined as follows: Category 1: Share of households with the lowest share of children per household, Category 2: Share of households with the share of children per household very far below the average, Category 3: Share of households with the share of children per household far below the average, Category 4: Share of households with the share of children per household below the average, Category 5: Share of households with the share of children per household slightly below the average, Category 6: Share of households with an average share of children per household, Category 7: Share of households with the share of children per household slightly above the average, Category 8: Share of households with the share of children per household above the average, Category 9: Share of households with the highest share of children per household

Table 2.A.5. Representativeness with Respect to Socioeconomic Status

Characteristic	Germany	Sample	p-value
Share of HHs falling into Status Category 1	11.110	5.624	0.000
Share of HHs falling into Status Category 2	11.110	7.480	0.000
Share of HHs falling into Status Category 3	11.110	9.533	0.000
Share of HHs falling into Status Category 4	11.110	11.363	0.290
Share of HHs falling into Status Category 5	11.110	12.991	0.000
Share of HHs falling into Status Category 6	11.110	13.664	0.000
Share of HHs falling into Status Category 7	11.110	13.677	0.000
Share of HHs falling into Status Category 8	11.110	13.496	0.000
Share of HHs falling into Status Category 9	11.120	12.171	0.077

Notes: This table shows the average share (in percent) of households falling into different socioeconomic status categories for Germany and for our sample. The averages are generated as weighted averages, with the weights for Germany being calculated as the zipcode population divided by the total population of households in Germany ($\approx$ 41 million). For the sample, the weights are calculated as the number of households living in a zipcode area divided by the total number of households in our sample. The p-values stem from two-sided t-tests using the weights described earlier. A household's status is defined as an index value based on education and income. The categories 1-9 are again defined as percentiles of the Germany-wide distribution of this index, explaining why the shares of households in Germany falling into each category are approximately the same. The higher the category, the "higher" the status.

Table 2.A.6. Education and Income by Socioeconomic Status

Status Category	Share with Secondary School	Share with A-Level	Share with University Education	Avg. Net Income per Month
1	12.71	5.86	4.37	1.667
2	12.31	6.42	4.88	2.009
3	13.80	8.24	8.64	2.271
4	11.80	8.07	7.38	2.469
5	12.68	11.84	10.79	2.636
6	11.20	11.46	12.71	2.811
7	9.72	13.38	14.38	3.009
8	9.73	14.01	13.30	3.022
9	6.05	20.71	23.55	3.630

Notes: This table shows the share of household heads with a secondary school education (*Hauptschule*), with an A-Level education (*Abitur*) and the share of household heads who went to university as well as average net income by socioeconomic status category. We clearly see that, the higher the education level and income, the higher the status category. This is a mechanical result from how the status variable and categories are generated. Source: *microm (2021)*

Appendix 2.B Summary Statistics

Table 2.B.1. Summary Statistics

	Mean	Min	Max	Observations
Hourly Price[a]	3.86	-10.711	23.205	5’583’500
Hourly Wind Production[b]	14.23	0.136	46.137	5’583’500
Daily Average Temperature[c]	11.08	-10.300	31.700	5’583’500
Sunshine Hours per Day	4.95	0.000	16.783	5’545’518
Daily Average Air Pressure[d]	984.14	815.190	1044.870	5’389’395

Notes: This table presents summary statistics of the (net) electricity price, wind production and weather controls. Note that for readability all weather controls are divided by 10 in our regressions. Here, we kept them as they are reported by the German Meteorological Service.
[a] Net price in Cent. The final price depends on taxes and local levies. On average, these amount to approximately 20 Cent.
[b] In GW
[c] In degree Celsius
[d] In hPA

Appendix 2.C Price Effects in Control Sample

2.C.1 OLS-Estimates of the Control Sample

In this section, we present the OLS estimates of the original regression based on Table 2.2 for the control sample. We subsequently add different levels of fixed effects to underline their importance by showing that we would indeed estimate a positive price sensitivity if we would not include them.

Table 2.C.1. Effect on Electricity Consumption: OLS Estimates, Control Sample

Dependent Variable	**Hourly Electricity Consumption**					
	(1)	(2)	(3)	(4)	(5)	(6)
EPEX Spot Price	.017***	.017***	.022***	.003***	-.002***	.002***
	(.001)	(.002)	(.001)	(.001)	(.001)	(.001)
Daily Temperature						-.012*
						(.006)
Hours of Sunshine						-.005*
						(.003)
Air Pressure						-.004**
						(.002)
Household fixed-effects	No	Yes	Yes	Yes	Yes	Yes
DOW FEs	No	No	Yes	Yes	Yes	Yes
Hour FEs	No	No	No	Yes	Yes	Yes
Sample Month fixed-effects	No	No	No	No	Yes	Yes
Local Weather Controls	No	No	No	No	No	Yes
R^2_{within}	0.0000	0.0000	0.0156	0.0759	0.0831	0.0825
$R^2_{between}$	0.0000	0.0000	-	-	-	0.0192
$R^2_{overall}$	0.0075	0.0075	0.0129	0.0625	0.0685	0.0703
Observations	1’235’130	1’235’130	1’235’130	1’235’130	1’235’130	1’102’116
Households	130	130	130	130	130	116

Notes: For readability, local weather controls (measured in degrees celsius, hours, and hPA, respectively) are divided by 10. That is, the effect on hourly electricity consumption of an increase in temperature by one degree celsius is $\hat{\beta}_{temp}/10$

Results stem from OLS regression of hourly electricity consumption on the hourly electricity price per kWh, based on equation (2.4), without instrumenting for the hourly electricity price. We subsequently add household-, day-of-the-week-, hour-of-the-day- and sample month fixed effects. Finally, we add local weather controls in column (6), based on the weather data provided by the closest weather station. Standard errors are clustered at the household level.

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

2.C.2 Price Effects Over the Day and Over the Week

In order to underline that we indeed do not observe any price effect in our control sample, we recreate the same regressions on which Figures 2.11 and 2.12 are based on in the control sample. As Figure 2.C.1 shows, we clearly see that the estimated price effect is not significant for any hour of the day and that the relative magnitude of the variation in the reduced form dominates the variation in the first stage in both specifications. We also do not find any significant price effects for any day of the week, as shown in Figure 2.C.2.

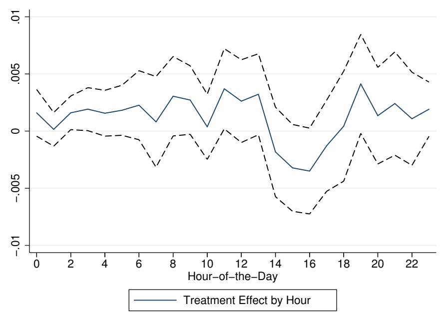

(a) Price Effect by Hour-of-the-Day in Control Sample

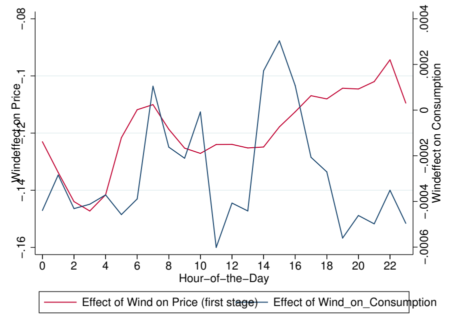

(b) Association of Wind Production and Electricity Consumption

Figure 2.C.1. Price Effect and Association of Wind Production and Electricity Consumption

Figure 2.C.1a presents the estimated price effect for each hour of the dayin the control sample. The dashed line denotes the 95%-confidence interval. At no hour of the day we estimate a significant price effect. At the same time, Figure 2.C.1b reveals that the first stage pattern found in the original sample is almost unchanged. The red line represents the first stage results over the day, measured on the left y-axis. The blue line represents the reduced form results, measured on the right y-axis. Standard errors are clustered at the household level.

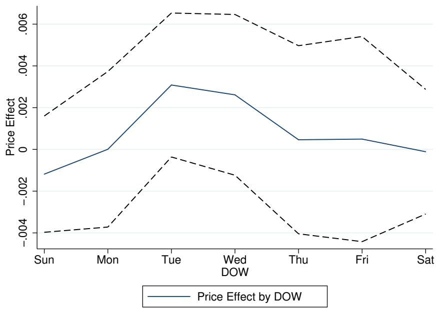

(a) TE by Day-of-the-Week

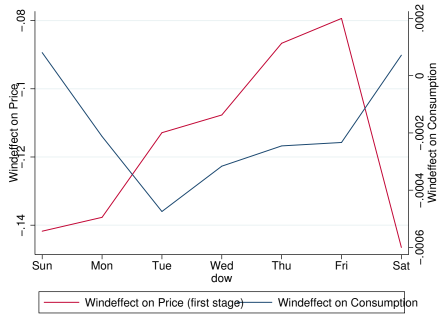

(b) Association of Wind Production and Electricity Consumption

Figure 2.C.2. Price Effect and Association of Wind Production and Electricity Consumption in Control Sample

Figure 2.C.2a presents the price effect for each day of the week in the control sample. The dashed line denotes the 95%-confidence interval. During no day of the week, we estimate a significant price effect. At the same time, Figure 2.C.2 reveals that the first stage pattern found in the original sample is almost unchanged compared to Figure 2.12a. As before, the red line represents the first stage results over the week, measured on the left y-axis. The blue line represents the reduced form results, measured on the right y-axis. Standard errors are clustered at the household level.

Appendix 2.D Robustness Checks

In order to investigate the robustness of our results, we conduct several sensitivity checks on different subsamples of our datasets. In particular, we will analyze whether our results hold up if we restrict our analysis to certain hours of the day and days of the week. This, along with Figures 2.11a and 2.12a ensures that our estimates are not driven by hours and days of particularly high price elasticity. Additionally, we will run separate analyzes on observations before and during the COVID-19 pandemic in order to analyze whether the increased necessity and possibility of working from home affected the price sensitivity. We also conduct our main analyses on different sets of household clusters according to our cluster analysis in Section 2.5.1. Finally, we also check whether our results hold up for different specifications and choice of instrument and clustering level.

2.D.1 Estimation Without the Nighttime Hours

Excluding the hours between 11 p.m. and 6 a.m. from the regression allows us to ensure that our findings are not driven by households using excessive amounts of energy during the cheap nighttime hours, for example by using an API. Even though we conduct a cluster analysis in Section 2.5.1 and use the Austrian dataset to show the load profiles of households using an API or an IM, excluding the nighttime hours from the analysis helps to ensure the validity of our results. Tables 2.D.1 and 2.D.2 show that our results hold up even if we exclude observations during the night from 10 p.m. to 6 a.m.

Table 2.D.1. Effect on Electricity Consumption: IV Estimates of Intra-Day Price Effects, excluding nighttime consumption

Dependent Variable	**Hourly Electricity Consumption**						
	(1)	(2)	(3)	(4)	(5)	(6)	(Control)
EPEX Spot Price	-.073***[a]	-.072***	-.066***	-.013***	-.028***	-.028***	.001
	(.000)	(.003)	(.003)	(.001)	(.002)	(.002)	(.002)
Daily Temperature						-.100***	.013
						(.019)	(.020)
Hours of Sunshine						-.080***	-.006
						(.006)	(.004)
Air Pressure						-.005	-.004
						(.008)	(.010)
Household fixed-effects	No	Yes	Yes	Yes	Yes	Yes	Yes
Hour FEs	No	No	Yes	Yes	Yes	Yes	Yes
Sample Month fixed-effects	No	No	No	Yes	Yes	Yes	Yes
Date fixed-effects	No	No	No	No	Yes	Yes	Yes
Local Weather Controls	No	No	No	No	No	Yes	Yes
$R^2_{centered}$	-	-0.0489	-0.0337	0.0016	0.0001	0.0013	-0.0000
$R^2_{uncentered}$	-	-0.0489	-0.0337	0.0016	0.0001	0.0013	-0.0000
Observations	3’957’396	3’957’396	3’957’396	3’957’396	3’957’396	3’799’268	780’912
Households	829	829	829	829	829	796	116

Notes: For readability, local weather controls (measured in degrees celsius, hours, and hPA, respectively) are divided by 10. That is, the effect on hourly electricity consumption of an increase in temperature by one degree celsius is $\hat{\beta}_{temp}/10$

Results stem from IV estimation of hourly electricity consumption on the hourly electricity price per kWh, based on equation (2.6), with the hourly eletricity price being instrumented for using hourly wind production in Germany according to equation (2.5). For this specification, we exluded observations during the nighttime hours between 11 p.m. and 6 a.m. We subsequently add household-, hour-of-the-day-, sample month- and date-fixed effects. Finally, we add local weather controls in column (6), based on the weather data provided by the closest weather station. In the last column, we run the same regression displayed in column (6) on the control sample. Excluding nighttime electricity consumption from the estimation (naturally) delivers (given that average hourly electricity consumption excluding nighttime consumption is 0.541 and the average hourly EPEX price exluding these hours is 24.2 ct/kWh) delivers an electricity price elasticity of -1.263 (compared to the calculated price elasticity of -1.668 for the whole day). Standard errors are clustered at the household level.

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

Table 2.D.2. Effect on Daily Electricity Consumption: IV Estimates, excluding nighttime consumption

Dependent Variable	**Daily Electricity Consumption**					
	(1)	(2)	(3)	(4)	(5)	(Control)
EPEX Spot Price	-1.227***	-1.226***	-1.273***	-.131***	-.155***	.023
	(.061)	(.061)	(.066)	(.023)	(.023)	(.020)
Daily Temperature					-1.349***	-.247**
					(.122)	(.098)
Hours of Sunshine					-1.130***	-.126**
					(.086)	(.052)
Air Pressure					-.067***	-.074**
					(.023)	(.030)
Household fixed-effects	No	Yes	Yes	Yes	Yes	Yes
DOW FEs	No	No	Yes	Yes	Yes	Yes
Sample Month fixed-effects	No	No	No	Yes	Yes	Yes
Local Weather Controls	No	No	No	No	Yes	Yes
R^2_{within}	0.0012	-	-	0.1388	0.1474	0.0722
$R^2_{between}$	0.0520	0.0520	0.0517	0.0606	0.0675	0.0139
$R^2_{overall}$	0.0021	0.0021	0.0042	0.0978	0.1023	0.0387
Observations	232’984	232’984	232’984	232’984	223’675	45’936
Households	829	829	829	829	796	116

Notes: For readability, local weather controls (measured in degrees celsius, hours, and hPA, respectively) are divided by 10. That is, the effect on daily electricity consumption of an increase in temperature by one degree celsius is $\hat{\beta}_{temp}/10$
Results stem from IV estimation of daily electricity consumption on the average hourly electricity price per kWh over the day, based on equation (2.8), with the average hourly eletricity price over the day being instrumented for using average hourly wind production over the day in Germany according to equation (2.7). We exclude observations during the nighttime hours between 22 p.m. and 6 a.m. for this analysis. We subsequently add household-, day-of-the-week- and sample month-fixed effects. Finally, we add local weather controls in column (5), based on the weather data provided by the closest weather station. In the last column, we run the same regression displayed in column (5) on the control sample. Given that average daily electricity consumption excluding nighttime consumption is 9.183 kWhs and the average hourly price is 24.2 ct/kWh, this delivers an electricity price elasticity of -0.408 (compared to the calculated inter-day price elasticity of -0.453 using observations over the whole day). Standard errors are clustered at the household level.
*** $p < 0.01$ ** $p < 0.05$ * $p < 0.1$

2.D.2 Interaction with COVID Pandemic

The COVID-19 pandemic affected many aspects of our life. In particular, due to lockdowns and social distancing, working from home became more and more viable over the course of 2020. As shown in Section 2.A, we find that the environment of the households in our sample indicates a higher socioeconomic status, which likely correlates with White Collar jobs that can be done from home. We now analyze how this increased availability of working from home affected households' price sensitivity. For this, we conduct our main analyses on two separate samples, consisting of observations before and after March 20, 2020.

Tables 2.D.3, 2.D.4, 2.D.5 and 2.D.6 show that intra-day demand elasticities increased slightly in magnitude over the course of the pandemic, whereas inter-day demand elasticities decreased slightly. This indicates that households were now able to shift electricity consumption during hours in which they usually were not at home, while not being able to shift electricity consumption across days as much as before. However, note that the absolute differences in the estimated price sensitivities are small.

Table 2.D.3. Effect on Electricity Consumption: IV Estimates of Intra-Day Price Effects, before Covid-Pandemic

Dependent Variable	**Hourly Electricity Consumption** (1)	(2)	(3)	(4)	(5)	(6)
EPEX Spot Price	-.062***[a]	-.052***	-.049***	-.011***	-.034***	-.035***
	(.001)	(.005)	(.004)	(.003)	(.004)	(.002)
Daily Temperature						-.099***
						(.024)
Hours of Sunshine						-.062***
						(.008)
Air Pressure						-.007
						(.011)
Household fixed-effects	No	Yes	Yes	Yes	Yes	Yes
Hour FEs	No	No	Yes	Yes	Yes	Yes
Sample Month fixed-effects	No	No	No	Yes	Yes	Yes
Date fixed-effects	No	No	No	No	Yes	Yes
Local Weather Controls	No	No	No	No	No	Yes
$R^2_{centered}$	-	-0.0033	0.0005	0.0010	0.0009	0.0016
$R^2_{uncentered}$	-	-0.0033	0.0005	0.0010	0.0009	0.0016
Observations	1’410’690	1’410’690	1’410’690	1’410’690	1’410’690	1’368’058
Households	377	377	377	377	377	368

Notes: For readability, local weather controls (measured in degrees celsius, hours, and hPA, respectively) are divided by 10. That is, the effect on hourly electricity consumption of an increase in temperature by one degree celsius is $\hat{\beta}_{temp}/10$
Results stem from IV estimation of hourly electricity consumption on the hourly electricity price per kWh, based on equation (2.6), with the hourly electricity price being instrumented for using hourly wind production in Germany according to equation (2.5). We subsequently add household-, hour-of-the-day-, sample month- and date-fixed effects. Finally, we add local weather controls in column (6), based on the weather data provided by the closest weather station. We only include observations up to March 1, 2020 for this analysis. As all observations in the control sample fall in this period, we exclude the control sample from this analysis. Given that average hourly electricity consumption in our sample before March 1, 2020 was 0.551 kWh and the average hourly EPEX price before March 1, 2020 was 24.0 ct/kWh, we calculate a price elasticity of -1.484. Standard errors are clustered at the household level.
*** $p < 0.01$ ** $p < 0.05$ * $p < 0.1$

Table 2.D.4. Effect on Electricity Consumption: IV Estimates of Intra-Day Price Effects, during Covid-Pandemic

Dependent Variable	**Hourly Electricity Consumption** (1)	(2)	(3)	(4)	(5)	(6)
EPEX Spot Price	-.072***[a]	-.064***	-.062***	-.016***	-.038***	-.038***
	(.001)	(.003)	(.003)	(.001)	(.002)	(.003)
Daily Temperature						-.103***
						(.020)
Hours of Sunshine						-.052***
						(.005)
Air Pressure						-.002
						(.008)
Household fixed-effects	No	Yes	Yes	Yes	Yes	Yes
Hour FEs	No	No	Yes	Yes	Yes	Yes
Sample Month fixed-effects	No	No	No	Yes	Yes	Yes
Date fixed-effects	No	No	No	No	Yes	Yes
Local Weather Controls	No	No	No	No	No	Yes
$R^2_{centered}$	-	-0.0372	-0.0261	0.0010	0.0002	0.0008
$R^2_{uncentered}$	-	-0.0372	-0.0261	0.0010	0.0002	0.0008
Observations	4’172’810	4’172’810	4’172’810	4’172’810	4’172’810	3’992’337
Households	820	820	820	820	820	787

Notes: For readability, local weather controls (measured in degrees celsius, hours, and hPA, respectively) are divided by 10. That is, the effect on hourly electricity consumption of an increase in temperature by one degree celsius is $\hat{\beta}_{temp}/10$
Results stem from IV estimation of hourly electricity consumption on the hourly electricity price per kWh, based on equation (2.6), with the hourly electricity price being instrumented for using hourly wind production in Germany according to equation (2.5). We subsequently add household-, hour-of-the-day-, sample month- and date-fixed effects. Finally, we add local weather controls in column (6), based on the weather data provided by the closest weather station. We exclude observations up to March 1, 2020 for this analysis. As all observations in the control sample fall in this period, we exclude the control sample from this analysis. Given that average hourly electricity consumption starting March 1, 2020 was 0.523 and the average hourly EPEX price after March 1, 2020 was 23.8 ct/kWh, we calculate a price elasticity of −1.729. Standard errors are clustered at the household level.
*** $p < 0.01$ ** $p < 0.05$ * $p < 0.1$

Table 2.D.5. Effect on Daily Electricity Consumption: IV Estimates, before Covid

Dependent Variable	**Daily Electricity Consumption**				
	(1)	(2)	(3)	(4)	(5)
EPEX Spot Price	-1.272***	-1.268***	-1.306***	-.090	-.299***
	(.117)	(.118)	(.125)	(.073)	(.069)
Daily Temperature					-1.489***
					(.255)
Hours of Sunshine					-1.203***
					(.167)
Air Pressure					-.009
					(.049)
Household fixed-effects	No	Yes	Yes	Yes	Yes
DOW FEs	No	No	Yes	Yes	Yes
Sample Month fixed-effects	No	No	No	Yes	Yes
Local Weather Controls	No	No	No	No	Yes
R^2_{within}	0.0098	-	-	0.1190	0.1205
$R^2_{between}$	0.0256	0.0256	0.0256	0.0509	0.0581
$R^2_{overall}$	0.0081	0.0081	0.0081	0.0739	0.0769
Observations	58’930	58’930	58’930	58’930	57’150
Households	377	377	377	377	368

Notes: For readability, local weather controls (measured in degrees celsius, hours, and hPA, respectively) are divided by 10. That is, the effect on daily electricity consumption of an increase in temperature by one degree celsius is $\hat{\beta}_{temp}/10$

Results stem from IV estimation of daily electricity consumption on the average hourly electricity price per kWh over the day, based on equation (2.8), with the average hourly electricity price over the day being instrumented for using average hourly wind production over the day in Germany according to equation (2.7). We subsequently add household-, day-of-the-week- and sample month-fixed effects. Finally, we add local weather controls in column (5), based on the weather data provided by the closest weather station. We only include observations up to March 1, 2020 for this analysis. As all observations in the control sample fall in this period, we exclude the control sample from this analysis. Given that average daily electricity consumption before March 1, 2020 was 13.184 kWhs and the average hourly price before March 1, 2020 was 24.0 ct/kWh, we calculate a price elasticity of -0.544. Standard errors are clustered at the household level.
*** $p < 0.01$ ** $p < 0.05$ * $p < 0.1$

Table 2.D.6. Effect on Daily Electricity Consumption: IV Estimates, during Covid

Dependent Variable	**Daily Electricity Consumption** (1)	(2)	(3)	(4)	(5)
EPEX Spot Price	-1.638***	-1.634***	-1.689***	-.219***	-.224***
	(.078)	(.078)	(.083)	(.035)	(.035)
Daily Temperature					-1.955***
					(.175)
Hours of Sunshine					-1.165***
					(.100)
Air Pressure					-.169***
					(.037)
Household fixed-effects	No	Yes	Yes	Yes	Yes
DOW FEs	No	No	Yes	Yes	Yes
Sample Month fixed-effects	No	No	No	Yes	Yes
Local Weather Controls	No	No	No	No	Yes
R^2_{within}	0.0060	-	-	0.1502	0.1594
$R^2_{between}$	0.0679	0.0679	0.0676	0.0644	0.0677
$R^2_{overall}$	0.0080	0.0080	0.0114	0.1071	0.1105
Observations	174’062	174’062	174’062	174’062	166’533
Households	820	820	820	820	787

Notes: For readability, local weather controls (measured in degrees celsius, hours, and hPA, respectively) are divided by 10. That is, the effect on daily electricity consumption of an increase in temperature by one degree celsius is $\hat{\beta}_{temp}/10$
Results stem from IV estimation of daily electricity consumption on the average hourly electricity price per kWh over the day, based on equation (2.8), with the average hourly electricity price over the day being instrumented for using average hourly wind production over the day in Germany according to equation (2.7). We subsequently add household-, day-of-the-week- and sample month-fixed effects. Finally, we add local weather controls in column (5), based on the weather data provided by the closest weather station. We exclude observations up to March 1, 2020 for this analysis. As all observations in the control sample fall in this period, we exclude the control sample from this analysis. Given that average daily electricity consumption starting on March 1, 2020 was 12.550 kWhs and the average hourly price starting on March 1, 2020 was 23.8 ct/kWh, we calculate a price elasticity of -0.423. Standard errors are clustered at the household level.
*** $p < 0.01$ ** $p < 0.05$ * $p < 0.1$

2.D.3 Weekends and Weekdays

We additionally conduct our analyses separately during weekends and weekdays. Of course, we expect the estimated demand elasticities to be significantly larger during the weekends, as people are at home on weekends, which not only drives up average consumption during the weekend, but also increases the scope for consumption shifting due to people being able to actively react to prices. As expected, we find significantly larger price elasticities (both intra- and inter-day) during the weekends than during weekdays.

Table 2.D.7. Effect on Electricity Consumption: IV Estimates of Intra-Day Price Effects on Weekends

Dependent Variable	**Hourly Electricity Consumption**						
	(1)	(2)	(3)	(4)	(5)	(6)	(Control)
EPEX Spot Price	-.073***[a]	-.073***	-.068***	-.023***	-.055***	-.054***	.001
	(.001)	(.004)	(.003)	(.002)	(.004)	(.004)	(.002)
Daily Temperature						-.105***	.008
						(.022)	(.017)
Hours of Sunshine						-.060***	-.001
						(.007)	(.007)
Air Pressure						-.010	-.004
						(.008)	(.010)
Household fixed-effects	No	Yes	Yes	Yes	Yes	Yes	Yes
Hour FEs	No	No	Yes	Yes	Yes	Yes	Yes
Sample Month fixed-effects	No	No	No	Yes	Yes	Yes	Yes
Date fixed-effects	No	No	No	No	Yes	Yes	Yes
Local Weather Controls	No	No	No	No	No	Yes	Yes
$R^2_{centered}$	-	-0.0244	-0.0158	0.0020	0.0019	0.0026	0.0001
$R^2_{uncentered}$	-	-0.0244	-0.0158	0.0020	0.0019	0.0026	0.0001
Observations	1’588’102	1’588’102	1’588’102	1’588’102	1’588’102	1’524’651	311’576
Households	829	829	829	829	829	796	116

Notes: For readability, local weather controls (measured in degrees celsius, hours, and hPA, respectively) are divided by 10. That is, the effect on hourly electricity consumption of an increase in temperature by one degree celsius is $\hat{\beta}_{temp}/10$
Results stem from IV estimation of hourly electricity consumption on the hourly electricity price per kWh, based on equation (2.6), with the hourly electricity price being instrumented for using hourly wind production in Germany according to equation (2.5). We subsequently add household-, hour-of-the-day-, sample month- and date-fixed effects. Finally, we add local weather controls in column (6), based on the weather data provided by the closest weather station. In the last column, we run the same regression displayed in column (6) on the control sample. We only include observations on Saturdays and Sundays for this analysis. Given that average hourly electricity consumption on weekends is 0.564 kWhs and the average EPEX price is 22.9 ct/kWh, we calculate a price elasticity of -2.193. Standard errors are clustered at the household level.
*** $p < 0.01$ ** $p < 0.05$ * $p < 0.1$

Table 2.D.8. Effect on Electricity Consumption: IV Estimates of Intra-Day Price Effects on Weekdays

Dependent Variable	**Hourly Electricity Consumption** (1)	(2)	(3)	(4)	(5)	(6)	(Control)
EPEX Spot Price	-.076***[a]	-.074***	-.072***	-.008***	-.025***	-.025***	.001
	(.001)	(.004)	(.004)	(.001)	(.002)	(.002)	(.001)
Daily Temperature						-.100***	.001
						(.019)	(.020)
Hours of Sunshine						-.052***	-.004
						(.006)	(.004)
Air Pressure						-.001	-.009
						(.009)	(.011)
Household fixed-effects	No	Yes	Yes	Yes	Yes	Yes	Yes
Hour FEs	No	No	Yes	Yes	Yes	Yes	Yes
Sample Month fixed-effects	No	No	No	Yes	Yes	Yes	Yes
Date fixed-effects	No	No	No	No	Yes	Yes	Yes
Local Weather Controls	No	No	No	No	No	Yes	Yes
$R^2_{centered}$	-	-0.0445	-0.0324	-0.0001	-0.0003	0.0004	0.0001
$R^2_{uncentered}$	-	-0.0445	-0.0324	-0.0001	-0.0003	0.0004	0.0001
Observations	3’995’398	3’995’398	3’995’398	3’995’398	3’995’398	3’835’744	790’540
Households	829	829	829	829	829	796	116

Notes: For readability, local weather controls (measured in degrees celsius, hours, and hPA, respectively) are divided by 10. That is, the effect on hourly electricity consumption of an increase in temperature by one degree celsius is $\hat{\beta}_{temp}/10$

Results stem from IV estimation of hourly electricity consumption on the hourly electricity price per kWh, based on equation (2.6), with the hourly electricity price being instrumented for using hourly wind production in Germany according to equation (2.5). We subsequently add household-, hour-of-the-day-, sample month- and date-fixed effects. Finally, we add local weather controls in column (6), based on the weather data provided by the closest weather station. In the last column, we run the same regression displayed in column (6) on the control sample. We exclude observations on Saturdays and Sundays for this analysis. Given that average hourly electricity consumption on weekdays is 0.517 kWhs and the average EPEX price is 24.3 ct/kWh, we calculate a price elasticity of -1.175. Standard errors are clustered at the household level.

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

Table 2.D.9. Effect on Daily Electricity Consumption: IV Estimates, Weekends

Dependent Variable	**Daily Electricity Consumption**					
	(1)	(2)	(3)	(4)	(5)	(Control)
EPEX Spot Price	-1.687***	-1.685***	-1.696***	-.354***	-.353***	.008
	(.089)	(.089)	(.090)	(.047)	(.049)	(.027)
Daily Temperature					-2.156***	-.491**
					(.205)	(.154)
Hours of Sunshine					-1.225***	-.122*
					(.140)	(.097)
Air Pressure					-.284***	-.162**
					(.051)	(.054)
Household fixed-effects	No	Yes	Yes	Yes	Yes	Yes
DOW FEs	No	No	Yes	Yes	Yes	Yes
Sample Month fixed-effects	No	No	No	Yes	Yes	Yes
Local Weather Controls	No	No	No	No	Yes	Yes
R^2_{within}	0.0006	-	-	0.1335	0.1419	0.0716
$R^2_{between}$	0.0441	0.0441	0.0441	0.0525	0.0510	0.0173
$R^2_{overall}$	0.0012	0.0012	0.0015	0.0918	0.0908	0.0385
Observations	66’231	66’231	66’231	66’231	63’585	12’992
Households	829	829	829	829	796	116

Notes: For readability, local weather controls (measured in degrees celsius, hours, and hPA, respectively) are divided by 10. That is, the effect on daily electricity consumption of an increase in temperature by one degree celsius is $\hat{\beta}_{temp}/10$
Results stem from IV estimation of daily electricity consumption on the average hourly electricity price per kWh over the day, based on equation (2.8), with the average hourly electricity price over the day being instrumented for using average hourly wind production over the day in Germany according to equation (2.7). We exclude observations during the nighttime hours between 22 p.m. and 6 a.m. for this analysis. We subsequently add household-, day-of-the-week- and sample month-fixed effects. Finally, we add local weather controls in column (5), based on the weather data provided by the closest weather station. In the last column, we run the same regression displayed in column (5) on the control sample. We only include observations on Saturdays and Sundays for this analysis. Given that average daily electricity consumption on weekends is 13.522 kWhs and the average hourly EPEX price is 22.9 ct/kWh, we calculate a price elasticity of −0.598. Standard errors are clustered at the household level.
*** $p < 0.01$ ** $p < 0.05$ * $p < 0.1$

Table 2.D.10. Effect on Daily Electricity Consumption: IV Estimates, Weekdays

Dependent Variable	**Daily Electricity Consumption** (1)	(2)	(3)	(4)	(5)	(Control)
EPEX Spot Price	-2.028***	-2.026***	-2.029***	-.082**	-.171***	.028
	(.110)	(.110)	(.111)	(.038)	(.036)	(.036)
Daily Temperature					-1.776***	-.275**
					(.171)	(.147)
Hours of Sunshine					-1.185***	-.088
					(.100)	(.066)
Air Pressure					-.074**	-.052
					(.035)	(.043)
Household fixed-effects	No	Yes	Yes	Yes	Yes	Yes
DOW FEs	No	No	Yes	Yes	Yes	Yes
Sample Month fixed-effects	No	No	No	Yes	Yes	Yes
Local Weather Controls	No	No	No	No	Yes	Yes
R^2_{within}	0.0109	-	-	0.1496	0.1567	0.0572
$R^2_{between}$	0.0439	0.0439	0.0439	0.0522	0.0588	0.0203
$R^2_{overall}$	0.0097	0.0097	0.0095	0.1006	0.1048	0.0273
Observations	166’761	166’761	166’761	166’761	160’098	32’944
Households	829	829	829	829	796	116

Notes: For readability, local weather controls (measured in degrees celsius, hours, and hPA, respectively) are divided by 10. That is, the effect on daily electricity consumption of an increase in temperature by one degree celsius is $\hat{\beta}_{temp}/10$

Results stem from IV estimation of daily electricity consumption on the average hourly electricity price per kWh over the day, based on equation (2.8), with the average hourly electricity price over the day being instrumented for using average hourly wind production over the day in Germany according to equation (2.7). We exclude observations during the nighttime hours between 22 p.m. and 6 a.m. for this analysis. We subsequently add household-, day-of-the-week- and sample month-fixed effects. Finally, we add local weather controls in column (5), based on the weather data provided by the closest weather station. In the last column, we run the same regression displayed in column (5) on the control sample. We exclude observations on Saturdays and Sundays for this analysis. Given that average daily electricity consumption on weekdays is 12.388 kWhs and the average hourly EPEX price is 24.3 ct/kWh, we calculate a price elasticity of −0.335. Standard errors are clustered at the household level.

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

2.D.4 Including All Four Clusters of Households

To show that excluding the high-consumption- and nighttime cluster from our analysis does not disproportionately distort our results, we also conduct our analyses on the whole sample, including households from all clusters. As Tables 2.D.11 and 2.D.12 show, including both clusters increases our estimates in magnitude. This result comes as no surprise, especially if we look at Figure 2.6. Households in the two clusters, which we originally excluded excessively using electricity during the night when it is cheap drives the negative estimate of the price sensitivity downwards. However, keep in mind that we purposely excluded both clusters from our analysis, as we believe the high-consumption cluster ($\approx$ 6% of our total sample) to consist of unrepresentative households and that the nighttime consumption cluster ($\approx$ 1.5%) does not necessarily allow for an analysis of behavioral mechanisms (for example, due to the usage of IMs).

Table 2.D.11. Effect on Electricity Consumption: IV Estimates of Intra-Day Price Effects, including all household clusters

Dependent Variable	**Hourly Electricity Consumption**					
	(1)	(2)	(3)	(4)	(5)	(6)
EPEX Spot Price	-.086***[a]	-.082***	-.079***	-.018***	-.046***	-.046***
	(.001)	(.004)	(.004)	(.002)	(.003)	(.003)
Daily Temperature						-.121***
						(.024)
Hours of Sunshine						-.061***
						(.006)
Air Pressure						-.004
						(.009)
Household fixed-effects	No	Yes	Yes	Yes	Yes	Yes
Hour FEs	No	No	Yes	Yes	Yes	Yes
Sample Month fixed-effects	No	No	No	Yes	Yes	Yes
Date fixed-effects	No	No	No	No	Yes	Yes
Local Weather Controls	No	No	No	No	No	Yes
$R^2_{centered}$	-	-0.0294	-0.0223	0.0010	0.0009	0.0015
$R^2_{uncentered}$	-	-0.0294	-0.0223	0.0010	0.0009	0.0015
Observations	6’004’525	6’004’525	6’004’525	6’004’525	6’004’525	5’758’445
Households	899	899	899	899	899	861

Notes: For readability, local weather controls (measured in degrees celsius, hours, and hPA, respectively) are divided by 10. That is, the effect on hourly electricity consumption of an increase in temperature by one degree celsius is $\hat{\beta}_{temp}/10$

Results stem from IV estimation of hourly electricity consumption on the hourly electricity price per kWh, based on equation (2.6), with the hourly electricity price being instrumented for using hourly wind production in Germany according to equation (2.5). We subsequently add household-, hour-of-the-day-, sample month- and date-fixed effects. Finally, we add local weather controls in column (6), based on the weather data provided by the closest weather station. We include all households in the sample for this analysis. That is, we also include households from the high consumption and nighttime consumption clusters displayed in Figure 2.6. Given that average hourly electricity consumption in our sample including all clusters is 0.591 kWhs and the average hourly EPEX price is 23.9 ct/kWh, we calculate a price elasticity of -1.860. Standard errors are clustered at the household level.

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

Table 2.D.12. Effect on Daily Electricity Consumption: IV Estimates, including all household clusters

Dependent Variable	**Daily Electricity Consumption** (1)	(2)	(3)	(4)	(5)
EPEX Spot Price	-2.083***	-2.082***	-2.171***	-.224**	-.300***
	(.106)	(.106)	(.114)	(.041)	(.038)
Daily Temperature					-2.283***
					(.234)
Hours of Sunshine					-1.312***
					(.118)
Air Pressure					-.170***
					(.035)
Household fixed-effects	No	Yes	Yes	Yes	Yes
DOW FEs	No	No	Yes	Yes	Yes
Sample Month fixed-effects	No	No	No	Yes	Yes
Local Weather Controls	No	No	No	No	Yes
R^2_{within}	0.0025	-	-	0.1312	0.1395
$R^2_{between}$	0.0416	0.0416	0.0412	0.0430	0.0493
$R^2_{overall}$	0.0030	0.0030	0.0048	0.0784	0.0879
Observations	250'567	250'567	250'567	250'567	240'298
Households	899	899	899	899	861

Notes: For readability, local weather controls (measured in degrees celsius, hours, and hPA, respectively) are divided by 10. That is, the effect on daily electricity consumption of an increase in temperature by one degree celsius is $\hat{\beta}_{temp}/10$

Results stem from IV estimation of daily electricity consumption on the average hourly electricity price per kWh over the day, based on equation (2.8), with the average hourly electricity price over the day being instrumented for using average hourly wind production over the day in Germany according to equation (2.7). We subsequently add household-, day-of-the-week- and sample month-fixed effects. Finally, we add local weather controls in column (5), based on the weather data provided by the closest weather station. We include all households in the sample for this analysis. That is, we also include households from the high consumption and nighttime consumption clusters displayed in Figure 2.6. Given that average daily electricity consumption in our sample including all clusters is 14.164 kWhs and the average hourly EPEX price is 23.9 ct/kWh, we calculate a price elasticity of -0.506. Standard errors are clustered at the household level.
*** $p < 0.01$ ** $p < 0.05$ * $p < 0.1$

2.D.5 Including only the Low-Consumption Cluster of Households

We also conduct our analysis on a sample only consisting of households belonging to the low consumption cluster. As Tables 2.D.13 and 2.D.14 show, we find that households are less price sensitive in absolute terms, but with lower average consumption in the low consumption cluster, the estimated price sensitivity translates to slightly stronger price elasticities, though the differences are small.

Table 2.D.13. Effect on Electricity Consumption: IV Estimates of Intra-Day Price Effects, including only the low consumption cluster

Dependent Variable	**Hourly Electricity Consumption** (1)	(2)	(3)	(4)	(5)	(6)
EPEX Spot Price	-.061***[a]	-.060***	-.059***	-.015***	-.030***	-.029***
	(.000)	(.004)	(.003)	(.002)	(.002)	(.003)
Daily Temperature						-.071***
						(.016)
Hours of Sunshine						-.047***
						(.004)
Air Pressure						-.001
						(.007)
Household fixed-effects	No	Yes	Yes	Yes	Yes	Yes
Hour FEs	No	No	Yes	Yes	Yes	Yes
Sample Month fixed-effects	No	No	No	Yes	Yes	Yes
Date fixed-effects	No	No	No	No	Yes	Yes
Local Weather Controls	No	No	No	No	No	Yes
$R^2_{centered}$	-	-0.0391	-0.0267	0.0016	0.0007	0.0013
$R^2_{uncentered}$	-	-0.0391	-0.0267	0.0016	0.0007	0.0013
Observations	4’285’562	4’285’562	4’285’562	4’285’562	4’285’562	4’141’446
Households	606	606	606	606	606	584

Notes: For readability, local weather controls (measured in degrees celsius, hours, and hPA, respectively) are divided by 10. That is, the effect on hourly electricity consumption of an increase in temperature by one degree celsius is $\hat{\beta}_{temp}/10$

Results stem from IV estimation of hourly electricity consumption on the hourly electricity price per kWh, based on equation (2.6), with the hourly electricity price being instrumented for using hourly wind production in Germany according to equation (2.5). We subsequently add household-, hour-of-the-day-, sample month- and date-fixed effects. Finally, we add local weather controls in column (6), based on the weather data provided by the closest weather station. We include only the “low consumption”-cluster for this analysis. Given that average hourly electricity consumption in the “low consumption”-cluster is 0.401 kWhs and the average hourly EPEX price is 23.9 ct/kWh, we calculate a price elasticity of -1.728. Standard errors are clustered at the household level.

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

Table 2.D.14. Effect on Daily Electricity Consumption: IV Estimates, including only the low consumption cluster

Dependent Variable	**Daily Electricity Consumption** (1)	(2)	(3)	(4)	(5)
EPEX Spot Price	-1.563***	-1.562***	-1.615***	-.224**	-.250***
	(.081)	(.081)	(.086)	(.041)	(.036)
Daily Temperature					-1.193***
					(.139)
Hours of Sunshine					-1.029***
					(.081)
Air Pressure					-.072**
					(.029)
Household fixed-effects	No	Yes	Yes	Yes	Yes
DOW FEs	No	No	Yes	Yes	Yes
Sample Month fixed-effects	No	No	No	Yes	Yes
Local Weather Controls	No	No	No	No	Yes
R^2_{within}	0.0016	-	-	0.1599	0.1657
$R^2_{between}$	0.0231	0.0231	0.0229	0.0364	0.0366
$R^2_{overall}$	0.0017	0.0017	0.0035	0.1247	0.1246
Observations	178'817	178'817	178'817	178'817	172'804
Households	606	606	606	606	861

Notes: For readability, local weather controls (measured in degrees celsius, hours, and hPA, respectively) are divided by 10. That is, the effect on daily electricity consumption of an increase in temperature by one degree celsius is $\hat{\beta}_{temp}/10$

Results stem from IV estimation of daily electricity consumption on the average hourly electricity price per kWh over the day, based on equation (2.8), with the average hourly electricity price over the day being instrumented for using average hourly wind production over the day in Germany according to equation (2.7). We subsequently add household-, day-of-the-week- and sample month-fixed effects. Finally, we add local weather controls in column (5), based on the weather data provided by the closest weather station. We include only the "low consumption" cluster for this analysis. Given that average daily electricity consumption for the low cluster is 9.610 kWhs and the average hourly EPEX price is 23.9 ct/kWh, we calculate a price elasticity of −0.622. Standard errors are clustered at the household level.

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

2.D.6 Using Wind Energy Production Forecast As the Instrument

As the day-ahead energy prices are generated based on wind production forecasts, it also helps to use wind production forecasts instead of actual wind production as the instrument in our specifications. With a correlation coefficient of 0.94 ($p=0.000$) we do not expect the corresponding results to differ much from our main results. Indeed, as Tables 2.D.15 and 2.D.16 confirm, our results are also robust to this alternative specification. The wind production forecast data used for this approach is provided by the information platform of the German transmission system operators (TSO)[23] and comprise day-ahead forecasts of the expected wind energy supply in the respective control area after direct marketing.

23. https://www.netztransparenz.de/Weitere-Veroeffentlichungen/Windenergie-Prognose, last visited Sep 27th, 2021

Table 2.D.15. Effect on Electricity Consumption: IV Estimates of Intra-Day Price Effects, using wind energy production forecast as the instrument

Dependent Variable	**Hourly Electricity Consumption**						
	(1)	(2)	(3)	(4)	(5)	(6)	(Control)
EPEX Spot Price	-.058***[a]	-.059***	-.057***	-.018***	-.033***	-.033***	.001
	(.000)	(.004)	(.003)	(.002)	(.003)	(.003)	(.001)
Daily Temperature						-.071***	.005
						(.016)	(.018)
Hours of Sunshine						-.047***	-.003
						(.004)	(.004)
Air Pressure						-.001	-.005
						(.007)	(.010)
Household fixed-effects	No	Yes	Yes	Yes	Yes	Yes	Yes
Hour FEs	No	No	Yes	Yes	Yes	Yes	Yes
Sample Month fixed-effects	No	No	No	Yes	Yes	Yes	Yes
Date fixed-effects	No	No	No	No	Yes	Yes	Yes
Local Weather Controls	No	No	No	No	No	Yes	Yes
$R^2_{centered}$	-	-0.0367	-0.0257	0.0015	0.0005	0.0011	0.0001
$R^2_{uncentered}$	-	-0.0367	-0.0257	0.0015	0.0005	0.0011	0.0001
Observations	4'285'562	4'285'562	4'285'562	4'285'562	4'285'562	4'141'446	1'102'116
Households	606	606	606	606	606	584	116

Notes: For readability, local weather controls (measured in degrees celsius, hours, and hPA, respectively) are divided by 10. That is, the effect on hourly electricity consumption of an increase in temperature by one degree celsius is $\hat{\beta}_{temp}/10$

Results stem from IV estimation of hourly electricity consumption on the hourly electricity price per kWh, based on equation (2.6), with the hourly electricity price being instrumented for using the hourly wind production forecast in Germany according to equation (2.5). We subsequently add household-, hour-of-the-day-, sample month- and date-fixed effects. Finally, we add local weather controls in column (6), based on the weather data provided by the closest weather station. In the last column, we run the same regression displayed in column (6) on the control sample. Given that average hourly electricity consumption is 0.530 kWhs and the average hourly EPEX price is 23.9 ct/kWh, we calculate a price elasticity of −1.488. Standard errors are clustered at the household level.

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

Table 2.D.16. Effect on Daily Electricity Consumption: IV Estimates, using daily average hourly wind energy production forecast as the instrument

Dependent Variable	**Daily Electricity Consumption**					
	(1)	(2)	(3)	(4)	(5)	(Control)
EPEX Spot Price	-1.522***	-1.522***	-1.562***	-.297**	-.292***	-.029
	(.081)	(.081)	(.085)	(.041)	(.040)	(.028)
Daily Temperature					-1.221***	-.309**
					(.139)	(.134)
Hours of Sunshine					-1.031***	-.100*
					(.081)	(.058)
Air Pressure					-.065**	-.086**
					(.029)	(.043)
Household fixed-effects	No	Yes	Yes	Yes	Yes	Yes
DOW FEs	No	No	Yes	Yes	Yes	Yes
Sample Month fixed-effects	No	No	No	Yes	Yes	Yes
Local Weather Controls	No	No	No	No	Yes	Yes
R^2_{within}	0.0016	-	-	0.1595	0.1655	0.1655
$R^2_{between}$	0.0231	0.0231	0.0229	0.0364	0.0367	0.0367
$R^2_{overall}$	0.0017	0.0017	0.0034	0.1244	0.1249	0.1249
Observations	178’817	178’817	178’817	178’817	172’804	172’804
Households	606	606	606	606	584	116

Notes: For readability, local weather controls (measured in degrees celsius, hours, and hPA, respectively) are divided by 10. That is, the effect on daily electricity consumption of an increase in temperature by one degree celsius is $\hat{\beta}_{temp}/10$

Results stem from IV estimation of daily electricity consumption on the average hourly electricity price per kWh over the day, based on equation (2.8), with the average hourly electricity price over the day being instrumented for using the average hourly wind production forecast over the day in Germany according to equation (2.7). We subsequently add household-, day-of-the-week- and sample month-fixed effects. Finally, we add local weather controls in column (5), based on the weather data provided by the closest weather station. In the last column, we run the same regression displayed in column (5) on the control sample. Given that average daily electricity consumption is 12.710 kWhs and the average hourly EPEX price is 23.9 ct/kWh, we calculate a price elasticity of −0.470. Standard errors are clustered at the household level.

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

2.D.7 Clustering Standard Errors at the Datetime- and Date Level

Table 2.D.17. Effect on Electricity Consumption: IV Estimates of Intra-Day Price Effects, S.E. clustered at the Date × Hour Level

Dependent Variable	**Hourly Electricity Consumption** (1)	(2)	(3)	(4)	(5)	(6)	(Control)
EPEX Spot Price	-.073***	-.072***	-.069***	-.015***	-.037***	-.037***	.001
	(.000)	(.003)	(.002)	(.001)	(.002)	(.003)	(.001)
Daily Temperature						-.101***	.005*
						(.003)	(.003)
Hours of Sunshine						-.055***	-.003*
						(.002)	(.002)
Air Pressure						-.005***	-.005***
						(.002)	(.001)
Household fixed-effects	No	Yes	Yes	Yes	Yes	Yes	Yes
Hour FEs	No	No	Yes	Yes	Yes	Yes	Yes
Sample Month fixed-effects	No	No	No	Yes	Yes	Yes	Yes
Date fixed-effects	No	No	No	No	Yes	Yes	Yes
Local Weather Controls	No	No	No	No	No	Yes	Yes
$R^2_{centered}$	-	-0.0362	-0.0249	0.0011	0.0006	0.0012	0.0001
$R^2_{uncentered}$	-	-0.0362	-0.0249	0.0011	0.0006	0.0012	0.0001
Observations	5’583’500	5’583’500	5’583’500	5’583’500	5’583’500	5’360’395	1’102’116
Households	829	829	829	829	829	796	116

Notes: For readability, local weather controls (measured in degrees celsius, hours, and hPA, respectively) are divided by 10. That is, the effect on hourly electricity consumption of an increase in temperature by one degree celsius is $\hat{\beta}_{temp}/10$ Results stem from IV estimation of hourly electricity consumption on the hourly electricity price per kWh, based on equation (2.6), with the hourly electricity price being instrumented for using the hourly wind production forecast in Germany according to equation (2.5). We subsequently add household-, hour-of-the-day-, sample month- and date-fixed effects. Finally, we add local weather controls in column (6), based on the weather data provided by the closest weather station. In the last column, we run the same regression displayed in column (6) on the control sample. Given that average hourly electricity consumption is 0.530 kWhs and the average hourly EPEX price is 23.9 ct/kWh, we calculate a price elasticity of −1.488. Standard errors are clustered at the date × hour level.
*** $p < 0.01$ ** $p < 0.05$ * $p < 0.1$

Table 2.D.18. Effect on Daily Electricity Consumption: IV Estimates, S.E. clustered at the Date × Time Level

Dependent Variable	**Daily Electricity Consumption**					
	(1)	(2)	(3)	(4)	(5)	(Control)
EPEX Spot Price	-1.858***	-1.837***	-1.903***	-.187**	-.241***	.026
	(.287)	(.253)	(.270)	(.074)	(.055)	(.016)
Daily Temperature					-1.855***	-.313***
					(.161)	(.067)
Hours of Sunshine					-1.195***	-.100*
					(.094)	(.053)
Air Pressure					-.133***	-.086***
					(.057)	(.021)
Household fixed-effects	No	Yes	Yes	Yes	Yes	Yes
DOW FEs	No	No	Yes	Yes	Yes	Yes
Sample Month fixed-effects	No	No	No	Yes	Yes	Yes
Local Weather Controls	No	No	No	No	Yes	Yes
$R^2_{centered}$	-0.0931	-0.1274	-0.1333	-0.0002	0.0092	0.0016
$R^2_{uncentered}$	0.5144	-0.1274	-0.1333	-0.0002	0.0092	0.0016
Observations	232’992	232’992	232’992	232’992	223’683	45’936
Households	829	829	829	829	796	116

Notes: For readability, local weather controls (measured in degrees celsius, hours, and hPA, respectively) are divided by 10. That is, the effect on daily electricity consumption of an increase in temperature by one degree celsius is $\hat{\beta}_{temp}/10$
Results stem from IV estimation of daily electricity consumption on the average hourly electricity price per kWh over the day, based on equation (2.8), with the average hourly electricity price over the day being instrumented for using the average hourly wind production forecast over the day in Germany according to equation (2.7). We subsequently add household-, day-of-the-week- and sample month-fixed effects. Finally, we add local weather controls in column (5), based on the weather data provided by the closest weather station. In the last column, we run the same regression displayed in column (5) on the control sample. Given that average daily electricity consumption is 12.710 kWhs and the average hourly EPEX price is 23.9 ct/kWh, we calculate a price elasticity of −0.470. Standard errors are clustered at the date level.
*** $p < 0.01$ ** $p < 0.05$ * $p < 0.1$

Appendix 2.E Differences as Outcome, Regressor and Instrument

Another way to estimate the daily price effect is to use a different outcome variable, regressor and instrument. That is, instead of regressing hourly electricity consumption on the electricity price, which is instrumented for with wind production, we regress the difference in hourly electricity consumption and the 24-hour lag in electricity consumption on the difference in the hourly electricity price and its 24-hour lag, which is instrumented with the difference in hourly wind production and its 24-hour lag. The advantage of this approach lies in its robustness to time trends. As Table 2.E.1 shows, time fixed effects do not change the results. We still estimate a significantly negative price sensitivity for households in our sample, whereas the estimates price sensitivity in the control sample is (though significant) small and positive.

Table 2.E.1. Effect on Electricity Consumption Difference: IV Estimates of Price Effects

Dependent Variable	**Hourly Electricity Consumption Difference to Day Before**						
	(1)	(2)	(3)	(4)	(5)	(6)	(Control)
EPEX Spot Price	-.029***[a]	-.029***	-.028***	-.028***	-.028***	-.028***	.002**
	(.002)	(.002)	(.002)	(.002)	(.002)	(.002)	(.001)
Daily Temperature						-.009***	-.000
						(.002)	(.000)
Hours of Sunshine						-.038***	-.000
						(.003)	(.000)
Air Pressure						-.004***	-.000
						(.001)	(.000)
Household fixed-effects	No	Yes	Yes	Yes	Yes	Yes	Yes
DOW FEs	No	No	Yes	Yes	Yes	Yes	Yes
Hour FEs	No	No	No	Yes	Yes	Yes	Yes
Sample Month fixed-effects	No	No	No	No	Yes	Yes	Yes
Local Weather Controls	No	No	No	No	No	Yes	Yes
R^2_{within}	0.0031	0.0031	0.0033	0.0033	0.0033	0.0035	0.0019
$R^2_{between}$	0.0452	0.0452	0.0268	0.0268	0.0456	0.0003	0.0019
$R^2_{overall}$	0.0032	0.0032	0.0033	0.0033	0.0033	0.0034	0.0019
Observations	5’561’865	5’561’865	5’561’865	5’561’865	5’561’865	5’339’665	1’098’984
Households	829	829	829	829	829	796	116

Notes: Note: For readability, local weather controls are divided by 10. That is, the effect on hourly electricity consumption of an increase in temperature by one degree celsius is $\hat{\beta}_{temp}/10$. Results stem from IV estimation. We use the difference in hourly electricity consumption and the 24-hour lag in electricity consumption as the outcome and the difference in hourly electricity price and its 24-hour lag, which is instrumented with the difference in hourly wind production and its 24-hour lag as the regressor. Standard errors are clustered at the household level.
*** $p < 0.01$ ** $p < 0.05$ * $p < 0.1$

Appendix 2.F Surplus Implications of Real-Time Pricing

While the advantages and possible drawbacks of dynamic electricity pricing for the energy market have been extensively discussed in Section 2.1, we also want to shed light on the implications on consumer surplus. For this exercise, we assume risk-neutral households and an electricity pricing scheme built around K electricity prices p_k, each being realized with probability q_k, such that, for the original price p_M, we have $p_M = \sum_k q_k p_k$. We can denote the expected change in consumer surplus from changing from a fixed-price to a dynamic pricing scheme as the sum of the consumer gains for each realized electricity price p_k, weighted by their respective probabilities q_k:

$$
\begin{aligned}
E[\Delta CS] = \sum_k \Delta CS_k &= -x_M \underbrace{\sum_k q_k \Delta p_k}_{=0 \text{ by construction}} - \frac{1}{2}\sum_k q_k \Delta p_k \underbrace{\Delta x_k}_{\beta \Delta p_k} \\
&= -\frac{1}{2}\beta \sum_k \Delta q_k p_k^2 \\
&= -\frac{1}{2}\underbrace{\beta}_{<0} V(p)
\end{aligned}
$$

Where x_M denotes electricity consumption in the case of the original price p_M, Δp_k denotes the difference $p_k - p_M$ and Δx_k denotes the difference $x_k - x_M$. Note that the expected change of the consumer surplus is always positive for a mean-preserving spread of the electricity prices. This result has a tangible intuition: For risk-neutral households, a fixed electricity price p_M results in the same surplus outcome as a mean-preserving spread of electricity prices while forcing households to always consume x_M, that is, forcing them to consume the same amount of energy they would have consumed if the price would have been fixed at $p_M = \sum_k q_k p_k$. However, with varying electricity prices, households can also adjust their electricity consumption, allowing them to be strictly better off as long as $p_k \neq p_M$ for at least one k.

To illustrate this result, consider the two-prices case with probability $q_k = 0.5$ and $\Delta p_H = -\Delta p_L$. Then:

$$
\begin{aligned}
\Delta CS &= -x_M \underbrace{\left(\frac{1}{2}\Delta p_H + \frac{1}{2}\Delta p_L\right)}_{=0} - \frac{1}{2}\left(\frac{1}{2}\Delta p_H \Delta x_H\right) \\
&= \frac{1}{2}\left(\frac{1}{2}\Delta p_H \Delta x_H + \frac{1}{2}\Delta p_L \Delta x_L\right)
\end{aligned}
$$

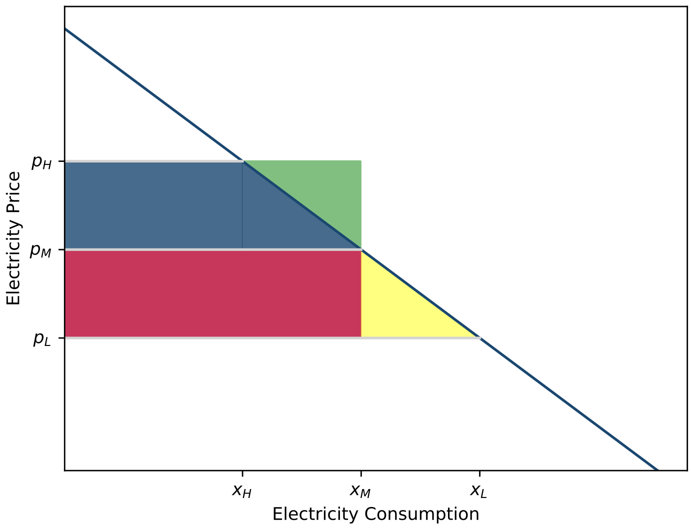

Figure 2.F.1. Daily Consumption Profile: Sample vs. Average German Household

This figure shows the effect of a high-low electricity pricing scheme on consumer surplus, assuming that prices increase and decrease with probability 0.5 by the same amount as compared to the original, fixed (= average) price. Note that x_L denotes consumption in case of a low price p_L, not low consumption per se. The same applies to x_H. The loss in consumer surplus in case of a high price is denoted by the dark blue trapezoid. The gain in consumer surplus in case of a low price is denoted by the red rectangle and the yellow triangle.

References

Alberini, Anna, and Massimo Filippini. 2011. “Response of Residential Electricity Demand to Price: The Effect of Measurement Error.” *Energy Economics* 33 (5): 889–95. DOI: https://doi.org/10.1016/j.eneco.2011.03.009. [72]

Allcott, Hunt. 2011. “Rethinking Real-Time Electricity Pricing.” *Resource and Energy Economics* 33 (4): 820–42. DOI: https://doi.org/10.1016/j.reseneeco.2011.06.003. [72, 104]

Andruszkiewicz, Jerzy, Józef Lorenc, and Agnieszka Weychan. 2019. “Demand Price Elasticity of Residential Electricity Consumers with Zonal Tariff Settlement Based on Their Load Profiles.” *Energies* 12 (22): 4317. DOI: 10.3390/en12224317. [72]

Beyertt, Adrian, Paul Verwiebe, Stephan Seim, Filip Milojkovic, and Joachim Müller-Kirchenbauer. 2020. “Felduntersuchung zu Behavioral Energy Efficiency Potentialen von privaten Haushalten.” *Zenondo*, DOI: 10.5281/zenodo.3855575. [76, 78]

Bollinger, Bryan K, and Wesley R Hartmann. 2020. “Information vs. Automation and Implications for Dynamic Pricing.” *Management Science* 66 (1): 290–314. DOI: https://doi.org/10.1287/mnsc.2018.3225. [72]

Boogen, Nina, Souvik Datta, and Massimo Filippini. 2014. “Estimating Residential Electricity Demand Using an Appliance Index and Energy Services.” ETH Zurich Working Paper. CER-ETH–Center of Economic Research, 200. DOI: http://dx.doi.org/10.2139/ssrn.2466015. [72]

Boogen, Nina, Souvik Datta, and Massimo Filippini. 2017. “Dynamic Models of Residential Electricity Demand: Evidence from Switzerland.” *Energy Strategy Reviews* 18: 85–92. DOI: https://doi.org/10.1016/j.esr.2017.09.010. [72]

Boonekamp, Piet G.M. 2007. “Price Elasticities, Policy Measures and Actual Developments in Household Energy Consumption – A Bottom Up Analysis for the Netherlands.” *Energy Economics* 29 (2): 133–57. DOI: https://doi.org/10.1016/j.eneco.2005.09.010. [72]

Borenstein, Severin. 2005. “The Long-Run Efficiency of Real-Time Electricity Pricing.” *Energy Journal* 26 (3): 93–116. DOI: 10.5547/ISSN0195-6574-EJ-Vol26-No3-5. [70]

Dahl, Gordon B., Andreas Ravndal Kostøl, and Magne Mogstad. 2014. “ Family Welfare Cultures.” *Quarterly Journal of Economics* 129 (4): 1711–52. DOI: 10.1093/qje/qju019. [90]

Degen, Kathrin, Charles Efferson, Fabian Frei, Lorenz Goette, and Raffael Lalive. 2013. “Smart Metering, Beratung oder Sozialer Vergleich - Was beeinflusst den Elektrizitätsverbrauch?” Report for the Swiss Federal Office of Energy. University of Lausanne. [76, 78, 79]

Fabra, Natalia, David Rapson, Mar Reguant, and Jingyuan Wang. 2021. “Estimating the Elasticity to Real-Time Pricing: Evidence from the Spanish Electricity Market.” *AEA Papers and Proceedings* 111: 425–29. DOI: 10.1257/pandp.20211007. [72, 86, 104]

Faruqui, Ahmad, and Stephen George. 2005. “Quantifying Customer Response to Dynamic Pricing.” *Electricity Journal* 18 (4): 53–63. DOI: https://doi.org/10.1016/j.tej.2005.04.005. [72]

Faruqui, Ahmad, Sanem Sergici, and Lamine Akaba. 2012. “Dynamic Pricing in a Moderate Climate: The Evidence from Connecticut.” *Available at SSRN 2028178*, DOI: http://dx.doi.org/10.2139/ssrn.2028178. [72]

Filippini, Massimo. 2011. “Short- and Long-run Time-of-Use Price Elasticities in Swiss Residential Electricity Demand.” *Energy Policy* 39 (10): 5811–17. DOI: https://doi.org/10.1016/j.enpol.2011.06.002. [72]

Forsa. 2015. “Akzeptanz von variablen Stromtarifen.” URL: https://www.vzbv.de/sites/default/files/downloads/Akzeptanz-variable-Stromtarife_Umfrage-Forsa-vzbv-November-2015.pdf (visited on 01/01/2022). [83, 103, 104]

Harding, Matthew, and Steven Sexton. 2017. "Household Response to Time-Varying Electricity Prices." *Annual Review of Resource Economics* 9 (1): 337–59. DOI: 10.1146/annurev-resource-100516-053437. [72]

Holtedahl, Pernille, and Frederick L. Joutz. 2004. "Residential Electricity Demand in Taiwan." *Energy Economics* 26 (2): 201–24. DOI: https://doi.org/10.1016/j.eneco.2003.11.001. [72]

Ito, Koichiro, Takanori Ida, and Makoto Tanaka. 2018. "Moral Suasion and Economic Incentives: Field Experimental Evidence from Energy Demand." *American Economic Journal: Economic Policy* 10 (1): 240–67. DOI: 10.1257/pol.20160093. [70, 72]

Jessoe, Katrina, and David Rapson. 2014. "Knowledge Is (Less) Power: Experimental Evidence from Residential Energy Use." *American Economic Review* 104 (4): 1417–38. DOI: 10.1257/aer.104.4.1417. [72]

Lijesen, Mark G. 2007. "The Real-Time Price Elasticity of Electricity." *Energy Economics* 29 (2): 249–58. DOI: https://doi.org/10.1016/j.eneco.2006.08.008. [72]

Okajima, Shigeharu, and Hiroko Okajima. 2013. "Estimation of Japanese Price Elasticities of Residential Electricity Demand, 1990–2007." *Energy Economics* 40: 433–40. DOI: https://doi.org/10.1016/j.eneco.2013.07.026. [72]

Patrick, Robert H, and Frank A Wolak. 2001. "Estimating the Customer-Level Demand for Electricity Under Real-Time Market Prices." Working Paper 8213. National Bureau of Economic Research. DOI: 10.3386/w8213. [72]

Reiss, Peter C., and Matthew W. White. 2005. "Household Electricity Demand, Revisited." *Review of Economic Studies* 72 (3): 853–83. DOI: 10.1111/0034-6527.00354. [72]

Schulte, Isabella, and Peter Heindl. 2017. "Price and Income Elasticities of Residential Energy Demand in Germany." *Energy Policy* 102: 512–28. DOI: https://doi.org/10.1016/j.enpol.2016.12.055. [72]

Silva, Susana, Isabel Soares, and Carlos Pinho. 2018. "Electricity Residential Demand Elasticities: Urban Versus Rural Areas in Portugal." *Energy* 144: 627–32. DOI: https://doi.org/10.1016/j.energy.2017.12.070. [72]

Trotta, Gianluca. 2020. "An Empirical Analysis of Domestic Electricity Load Profiles: Who Consumes How Much and When?" *Applied Energy* 275: 115399. DOI: https://doi.org/10.1016/j.apenergy.2020.115399. [83, 84]

Volland, Benjamin, and Ivan Tilov. 2018. "Price Elasticities of Electricity Demand in Switzerland: Results From a Household Panel." eng. IRENE Working Paper 18-03. Neuchâtel. URL: http://hdl.handle.net/10419/191512. [72]

Williams, Kevin R. 2021. "The Welfare Effects of Dynamic Pricing: Evidence from Airline Markets." Working Paper 28989. National Bureau of Economic Research. DOI: 10.3386/w28989. [104]

Wolak, Frank A. 2011. "Do Residential Customers Respond to Hourly Prices? Evidence from a Dynamic Pricing Experiment." *American Economic Review* 101 (3): 83–87. DOI: 10.1257/aer.101.3.83. [104]

Chapter 3

Job Search Autonomy

Joint with Patrick Arni and Amelie Schiprowski

3.1 Introduction

Modern unemployment policies tend to follow a paternalistic approach when it comes to the provision of job search effort. In many countries, Public Employment Services (PES) restrict the choice of effort through minimum application requirements and mandatory vacancy referrals. Such restrictions on the decision of how much and how broadly to search are commonly motivated by the concern that job seekers might under-provide search effort, for instance due to the disincentive effects of unemployment insurance (e.g., Schmieder, Wachter, and Bender (2016)) or behavioral phenomena, such as over-optimism (e.g. Spinnewijn (2015); Mueller, Spinnewijn, and Topa (2021)) or hyperbolic time preferences (e.g. DellaVigna and Paserman (2005)).

The degree of autonomy left to job seekers is a controversial policy choice. On the one hand, search effort is a key input to the process that matches unemployed individuals to jobs. Therefore, enforcing high effort may successfully foster labor market re-integration. At the same time, restrictions on search effort carry the risk of reducing intrinsic motivation and self-efficacy.[1] Moreover, effort restrictions involve an important general equilibrium dimension. Making all job seekers in a given labor market exert high effort will affect labor market tightness if job creation does not fully adjust. Depending on the relative importance of changes in individual effort versus changes in labor market tightness, it is unclear how the large-scale enforcement of search effort ultimately translates into job finding. Finally, there potentially exists a trade-off between the speed of job finding and job quality. If job seekers use

1. For evidence on intrinsic motivation, self-efficacy and self-regulation in the job search process, see Bandura (1977), Zimmerman, Boswell, Shipp, Dunford, and Boudreau (2012) and Guan, Deng, Sun, Wang, Cai, et al. (2013).

their autonomy to successfully direct their search towards higher paying jobs, this can have meaningful consequences for the individual and fiscal trade-offs involved.

In this paper, we provide comprehensive empirical evidence on the labor market effects of a large-scale policy change which reduced job search restrictions. The policy change *BernTop!* was implemented over the year 2012 in the Swiss canton Bern. It had the declared goal to promote the autonomy and self-efficacy of unemployed job seekers. Moreover, it aimed at reducing the administrative burden and improving the image of the PES. This resulted in a substantial reduction in the number of required job applications that job seekers faced. Job search requirements decreased by about 25% on average and mandatory vacancy referrals were almost completely abolished. These changes translated into a roughly proportional decrease in provided job applications. Moreover, job seekers reduced the occupational broadness of their job search, focusing on a smaller set of different occupations.

We estimate the reform's effect on labor market outcomes based on individual-level data from Swiss unemployment insurance and social security records. We set up a difference-in-differences framework, using job seekers registered in the rest of Switzerland as the control group. Before the reform, average labor market outcomes of job seekers inside and outside Bern evolved in a parallel way.

We first analyze how the average duration of unemployment spells was affected. We find that the reform-induced reduction in effort comes along with an increase in the length of unemployment spells by about 14 days (8%) on average. To fix ideas about the forces underlying this average effect, we discuss a simple conceptual framework. The framework illustrates the decomposition of the overall reform effect into a behavioral effect due to the decrease in individual effort and a tightness effect due to the decrease in aggregate effort. Moreover, the framework yields implications regarding the expected reform effect under different circumstances, which we assess with analyses of heterogeneity. First, we expect a higher decrease in job finding among job seekers whose individual effort decreased more strongly due to the reform. This conjecture is strongly supported in the data, where we observe that effects are most pronounced among job seekers whose effort is predicted to decrease most strongly due to the reform. Second, we expect job finding to react more in local labor markets with a higher initial labor market tightness. In the data, we we find suggestive evidence of stronger effects in labor markets with a high pre-reform vacancy-to-unemployed ratio. Third, we expect job finding to decrease more in local labor markets where the expected change in tightness due to the reform is lower. In line with this notion, we observe that the effect is strongest in labor markets with low commuting time to one of the adjacent cantons. In these markets, tightness is expected to decrease less, due to a higher exposure to commuters from outside Bern, whose search effort remains unchanged after the reform. We also find suggestive evidence that job seekers at the other side of the border benefited from the decrease in effort provided by job seekers in Bern and exited unemployment at a faster rate.

Taken together, the evidence suggests that tightness effects play a non-negligible role for the net effect of policies that target search effort.

We then study whether the increase in autonomy came at the benefit of higher re-employment earnings. In a model with directed search (e.g. Nekoei and Weber, 2017), the reform's effect on earnings is ambiguous: an increased duration of unemployment may reflect that job seekers become more selective about the jobs they apply to. This would on average result in higher post-unemployment earnings. On the other hand, a longer unemployment duration can lower average wage offers, for instance due to negative signaling. Our results suggest that the first effect dominates, as we estimate that the policy change increased average earnings by about 2.5 log points on average.

In a final step, we use our estimates to discuss the monetary trade-offs related to the reform's average effects. The aim is to quantify how the individual and fiscal costs of longer unemployment spells compare to the benefits of higher post-unemployment earnings. A simple back-of-the envelope calculation suggests that the earnings gains need to persist for about 6-7 years to offset the fiscal costs of longer unemployment spells. In turn, the worker's individual earnings losses due to longer unemployment are amortized after about 5 months if earnings gains persist.

This study relates to the literature on the effects of unemployment policies. A large body of literature has shown the positive relationship between UI generosity and unemployment duration (e.g., Card and Levine, 2000; Chetty, 2008; Lalive, 2008; Schmieder, Wachter, and Bender, 2012). This empirical relationship usually motivates the control which Public Employment Services (PES) exert over job seekers' effort in most OECD countries (Venn (2012)). With job search requirements and monitoring being a central measure in the toolbox of the PES, their impact has been studied in a multitude of contexts, with results mostly pointing towards a reduced length of unemployment spells (e.g., Johnson and Klepinger, 1994; Meyer, 1995; Klepinger, Johnson, and Joesch, 2002; Ashenfelter, Ashmore, and Deschênes, 2005; Van den Berg and Van der Klaauw, 2006; McVicar, 2008; Manning, 2009; McVicar, 2010; Van den Berg and Vikström, 2014; Arni and Schiprowski, 2019). We contribute with the comprehensive analysis of a large-scale policy change that drastically increased autonomy for all job seekers in a large, well-defined area. This distinguishes our setting from most previous studies, which typically rely on sources of variation that affect only a subset of job seekers in a given labor market. As real-world labor market reforms typically apply to the vast majority of job seekers in a given market, it is key to understand the effect of such large-scale changes, which affect not only individual behavior, but also labor market tightness. As our results show, these tightness effects matter for the outcomes of treated and untreated job seekers competing in the same labor market. As a result, the effect of changes in search effort is likely to be over-estimated when treatment and control group search in the same market.

By accounting for the relevance of tightness effects, we further relate to recent studies which acknowledge and estimate job search externalities. In the context of UI, market externalities have been shown to have an important impact on the optimal level of UI (Landais, Michaillat, and Saez, 2010; Landais, Michaillat, and Saez, 2018). Recent empirical evidence addresses this notion by documenting varying externalities of UI extensions (Lalive, Landais, and Zweimüller, 2015; Johnston and Mas, 2018) and job search assistance (Crépon, Duflo, Gurgand, Rathelot, and Zamora, 2013; Gautier, Muller, Klaauw, Rosholm, and Svarer, 2018; Cheung, Egebark, Forslund, Laun, Rodin, et al., 2019). While our study does not have the goal to separately identify the size of job search externalities, it shows that the local scope for search externalities through changes in labor market tightness is decisive for the average effect of large-scale changes in the search intensity of job seekers.

We proceed as follows: In Section 3.2, we describe the institutional setting and data, Section 3.3 describes the policy change and Section 3.4 presents the empirical design and discusses the results. Section 3.5 discusses the reform's monetary trade-offs and Section 3.6 concludes.

3.2 Data and Institutional Background

3.2.1 Data

Our empirical analysis is based on individual-level data from the Swiss UI registers provided by the Swiss State Secretariat for Economic Affairs (SECO), merged to social security records. The sample covers all unemployment spells starting between 2009 and 2015 of job seekers aged between 28 and 60. Additionally, we use data on the monthly job requirements set by the caseworkers for the cantons Bern, Fribourg, Solothurn and Tessin. These cantons cover about 22% of all UI recipients and three different geographic and language regions in Switzerland (Arni and Schiprowski, 2019). Overall, the dataset covers the date of registration for unemployment, the date and reason for deregistration from unemployment, the number of monthly application requirements in the mentioned cantons, the number and date of official referrals for each spell as well as rich information on socio-demographics and municipality level geolocation data of the place of residence of each job seeker in the dataset.[2] Moreover, we observe the history of pre- and post-unemployment earnings up to the end of 2015.

2. Note that the number of applications reported does not necessarily perfectly reflect actual job search effort. If the marginal costs for job seekers to report their search effort (i.e. reporting one additional application) are significant, then job seekers may only report as many applications as they have to, even if they sent out more. However, we assume that the marginal costs of reporting one additional application are sufficiently low such that job seekers also report applications that exceed the application requirements. This is also supported by our observation that on average, job seekers report more applications than required.

3.2.2 The Swiss Unemployment Insurance

In Switzerland, unemployed individuals are entitled to unemployment benefits if they contributed for at least twelve months during the two years prior to unemployment. To be eligible for the full benefit period, the contribution period extends up to 18 months for job seekers up to 55. Usually, the maximum potential benefits duration is 1.5 years for prime age workers, with variation with respect to the job seeker's employment history, age, and family situation. The replacement rate ranges between 70% and 80% of gross previous earnings, depending on the job seeker's family situation.

The process to claim unemployment benefits is strictly organized. As soon as an individual knows about her (upcoming) unemployment, she registers at the local Public Employment Service (PES) office, called the *Regionale Arbeitsvermittlung* (RAV). After registration, job seekers are assigned to a caseworker. Through regular meetings, caseworkers provide advice and counseling in the search process. Caseworkers also set application requirements and refer job seekers to vacancies according to general guidelines, which are set at the canton level.

3.2.3 The Application Requirements

The first caseworker meeting usually takes place around two to three weeks after registration. During this meeting, the caseworker sets the first application requirement, that is, the minimum number of monthly job applications which the job seeker must submit to avoid benefit cuts. Job seekers document their application activity in a monthly "protocol of search effort", which includes all types of applications made. The protocols are submitted on a monthly basis to the canton or to the PES office (depending on the canton), where they are collected and registered centrally. Job seekers are required to send in copies of their applications together with the protocols. Upon receiving the protocol, cantons or PES offices record the total number of applications in the central database. Caseworkers are legally obliged to assess whether the provided number of applications satisfies the requirement. They also check whether a minimum quality standard is met. Moreover, caseworkers occasionally verify the truthfulness of reported applications by calling the prospective employer. Once non-compliance with the search requirement is detected, a sanction can be imposed if the job seeker had no special reason or circumstance justifying the non-compliance.

3.2.4 The Referral Process

Caseworkers can officially refer job seekers to job openings if they believe to have found a fitting match in the PES database. These official referrals are, once made, mandatory to apply to and consist of several forms to be filled out both by the job seeker *and* the potential employer.

Among policy makers in Bern, the practice of requiring potential employers to give feedback on the applications was perceived as generating a burden to potential employers and leading to fewer employers actually reporting job openings to the PES database in order to avoid the additional administrative burden. Moreover, the vacancy referrals were considered to worsen the job prospects of job seekers through negative signaling effects.

3.3 The *BernTop!* Policy Change

Over the course of 2012, the department for economic affairs of the canton Bern enacted a policy change, *BernTop!*, which changed the strategy of the PES. The two main goals were to promote the autonomy and to increase the attractiveness of job seekers for employers, by improving the transparency in the job search process and minimizing "demotivators" in the job search process. Application requirements and vacancy referrals were regarded as the two major demotivators. Their use was decreased substantially.

- One of the most frequently mentioned demotivators and a controversial topic in the PES policy overall were the application requirements. In the official guide to the policy change, the PES notes that "We do not see the expedience of canton-wide application requirements. They do not improve the attractiveness of job seekers for potential employers, do not improve the PES' image and are inefficient from an administrative point-of-view." Over the course of *BernTop!*, the application requirements were reduced by approximately 25 %.
- The second "demotivator" tackled was the use of mandatory vacancy referrals. Through *BernTop!*, referrals were almost completely abolished.
- The number and definition of occupations a job seeker declared to search in was changed. Before *BernTop!*, job seekers had a vague catalogue of potential occupations to choose from. Over *BernTop!*, this catalogue was trimmed and job seekers were asked to only fill in occupations they saw as fitting and realistic to achieve.
- The time frame for the policy change was from August to December 2012. During this time, the PES offices trained their caseworkers to adapt the new policy.

3.3.1 Effects of the Policy Change in the Data

The reform *BernTop!* affected both ongoing and incoming unemployment spells. As a result, job seekers entering in 2013 were fully treated and job seekers entering in the previous months were partially treated. We therefore expect the change in requirements and vacancy referrals to become gradually visible. Panel (a) of Figure 3.1 is in line with this notion. During the years 2010 and 2011, the average requirement amounted to about 9 applications per month in Bern. Over the course of 2012,

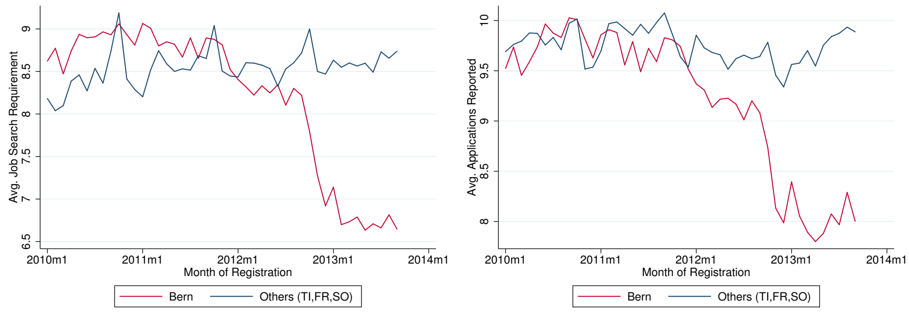

(a) No of Required Monthly Applications **(b)** No of Reported Monthly Applications

Figure 3.1. Required and Reported Applications

This figure shows the evolution of average required and reported applications in Bern and the other cantons. The x-axis denotes the month of registration for unemployment. As not all cantons report requirements in the data, the figure only includes Bern, Tessin, Fribourg and Solothurn. Note that *BernTop!* affected both ongoing and incoming spells, explaining the decrease in applications for spells registered in 2012.

this number fell down to about 6.5 applications per month for job seekers registered in 2013 or later. Panel (b) reveals that the drop in average application requirements translated to a comparable drop in average reported applications. Over the same time period, requirements and reported applications stayed roughly constant in the other cantons. Appendix Figure 3.B.1 additionally shows how the distribution of requirement and provided applications changed after the reform.

Figure 3.2a plots the share of unemployment spells that received at least one vacancy referral for Bern and the remaining cantons by month of registration. While this share stayed constant in the remaining cantons, it dropped from about 50% down to about 5% of spells over the course of the reform implementation (again, note that also spells registered prior to *BernTop!* are affected).

Finally, Figure 3.2b shows that the number of different occupations in which a job seeker states to search in decreased with the introduction of *BernTop!*. As a result, the self-reported occupational broadness of search also decreased, as shown in Appendix Figure 3.B.3.

In sum, we can see that as an immediate result of the reform, job search intensity decreased, vacancy referrals were stopped and job seekers started searching in fewer, more similar occupations.

3.4 Empirical Analysis

We aim at providing a comprehensive evaluation of the reform on the outcomes of unemployed job seekers. We consider two main aspects: (1) the effect on the duration of unemployment spells; and (2) the effect on post-unemployment earnings.

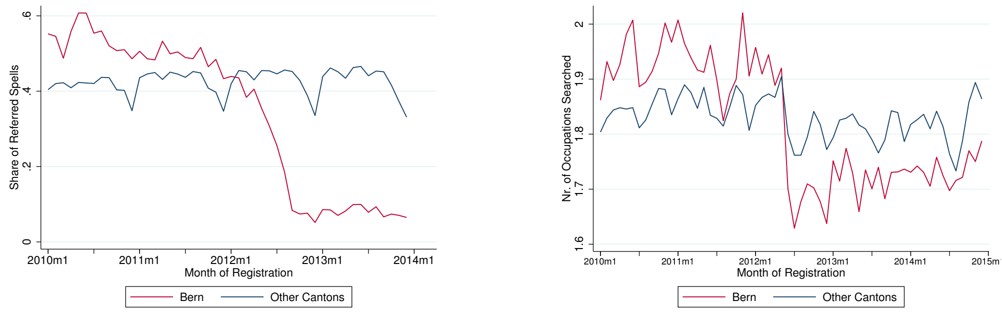

(a) Share of Spells That Were Referred at Least Once

(b) No of Occupations Searched

Figure 3.2. Measures of the Policy Change

Panel 3.2a shows the share of unemployment spells that received at least one referral. Panel 3.2b shows the average number of occupations job seekers declared to search in during their first meeting with a case-worker. Note that *BernTop!* affected both ongoing and incoming spells, explaining the decrease in referrals for spells registered in 2012.

3.4.1 Empirical Framework

We estimate the effect of the policy change on labor market outcomes using a difference-in-differences framework. The control group consists of all unemployment spells in Switzerland starting between January 2009 and December 2015 that were not located in Bern. We estimate the following dynamic difference-in-differences specification:

$$Y_i = \sum_{s=2009}^{2015} \gamma_s^{Bern} \mathbf{I}_{(y=s \,\&\, k=Bern)} + \delta \mathbf{I}_{(k=Bern)} + \tau_t + X_i' \beta + \epsilon_i \qquad (3.1)$$

Y_i describes a given labor market outcome of job seeker i. Indicators $I_{(y=s \,\&\, k=Bern)}$ equal one when a spell started in year s in the canton of Bern. As the reform's roll-out started in 2012, we use the year 2011 as the omitted baseline period. The indicator $\mathbf{I}_{(k=Bern)}$ equals one for all spells started in the canton of Bern and controls for time-constant differences between Bern and the other cantons. τ_t includes calendar month fixed effects, which control for aggregate time shocks in Switzerland. Controls for job seeker covariates are included in x_i (see summary statistics in Table 3.C.1 in Appendix 3.C). The difference-in-differences coefficients γ_s^{Bern} measure how outcomes in Bern changed compared to 2011, relative to the control group. Their causal interpretation for post-reform years relies on the key assumption that outcomes in Bern would on average have evolved in parallel to those in the other Swiss cantons. The estimates of γ_s^{Bern} for s=2009 and s=2010 inform about the relevance of this assumption during the pre-reform period.

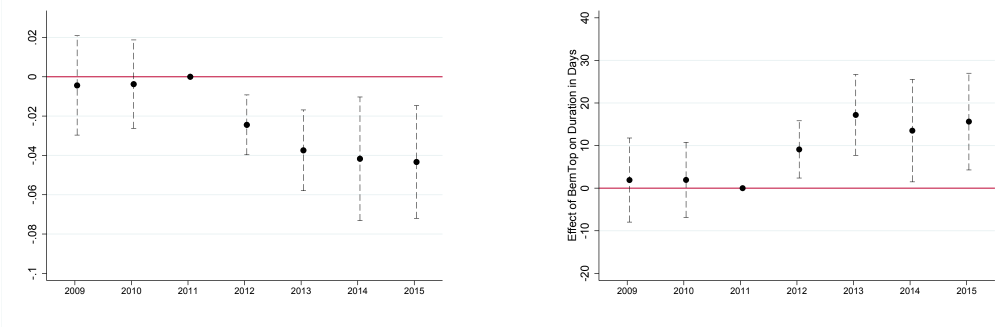

(a) Unemployment Exit Probability (< 12 months)

(b) Unemployment Duration in Days

Figure 3.3. Effect on Exit from Unemployment and Unemployment Duration

Estimated coefficients of a dynamic Diff-in-Diff framework (equation (3.1)). In panel (a), the outcome is the probability to exit unemployment within 12 months after entry. In panel (b), the outcome is the duration of unemployment (top-coded at 520 days). The treatment group includes unemployment spells registered in Bern, the control group includes unemployment spells registered in the rest of Switzerland. Regressions include controls for job seeker characteristics. The dashed lines denote 95% confidence intervals. Standard errors are clustered at the PES level (N=150).

We cluster standard errors at the level of the Public Employment Service (PES) office, which is the level at which requirement and referral policies are implemented.

3.4.2 Effect on Average Duration of Unemployment

Figure 3.3 shows the estimates of $\hat{\gamma}_s^{Bern}$. The baseline period is the pre-reform year 2011. Panel (a) reports effects on the probability to exit unemployment within 12 months after entry and panel (b) reports effects on the average duration of unemployment (top-coded at 520 days). Results show that the overall effect of the policy change on unemployment exit is negative. Among partially treated spells that started in 2012, the probability to exit within 12 months decreases by roughly two percentage points. The effect amounts to about four percentage points for spells starting in 2013 or later, after the reform was fully implemented. Accordingly, the duration of unemployment prolongs by about ten days on average for job seekers registered in 2012 and by about 12-18 days for job seekers registered thereafter.

Table 3.1 shows the corresponding pooled difference-in-differences estimates, excluding spells in 2012, which were partially treated. On average, the policy change prolonged the duration of unemployment by about 12.2 (columns 1, without controls) to 14.5 days (column 2, with controls). Compared to the pre-reform mean in Bern, this corresponds to an increase by 8%. As shown by columns (3) to (4), the effect seems to operate mostly during the first year of unemployment, although it remains negative and significant even when considering a time window of 18 months. Appendix Table 3.E.1 shows that the results are robust to various levels of clustering and fixed effects as well as different specifications and the inclusion of unemploy-

Table 3.1. Effect on Unemployment Duration and Exit: Pooled Diff-in-Diff Estimates

	UE Duration		P(Exit, 6 mon.)	P(Exit, 12 mon.)	P(Exit, 18 mon.)
	(1)	(2)	(3)	(4)	(5)
DiD	12.231***	14.472***	-0.033***	-0.038***	-0.027***
	(4.679)	(3.728)	(0.010)	(0.009)	(0.009)
Controls	No	Yes	Yes	Yes	Yes
Outcome Mean	266.27	266.27	0.45	0.66	0.80
R^2	0.026	0.114	0.093	0.086	0.094
N	348'448	348'448	348'448	348'448	348'448

This table reports estimates from a pooled version of the Difference-in-Differences framework defined by equation (3.1), excluding job seekers who registered in the year 2012. The treatment group includes unemployment spells registered in Bern, the control group includes unemployment spells registered in the rest of Switzerland. In columns (1) and (2), the outcome is the average unemployment duration in days (top-coded at 520 days). In columns (3) to (6), outcomes are the probability to exit unemployment after 6, 12, or 18 months, respectively. Summary statistics on control variables are reported in Appendix 3.C. Standard errors are clustered at the PES level (N=150). *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

ment spells starting in 2012. Anticipating the possibility of search externalities on job seekers registered outside of Bern, a further robustness check excludes from the control group job seekers in municipalities within less than 40 minutes commuting time to Bern. We observe only a small decrease in the results. Finally, Appendix Table 3.G.1 shows the effects on different sub-groups of job seekers. We observe particularly strong effects on the average unemployment duration of male job seekers, job seekers in Blue Collar occupations, low-skilled job seekers and job seekers without a Swiss nationality.

3.4.3 Role of Individual Effort and Labor Market Tightness

So far, the estimates have shown how the reform affected the average unemployment duration of job seekers in Bern. The size of this overall effect is likely determined two main forces: a decrease in individual-level effort provision and a change in labor market tightness due to the fact that job seekers in Bern collectively decrease their effort. In the following, we provide heterogeneity analyses that inform about the importance of these two counteracting forces.

To fix ideas, Appendix 3.A presents a simple conceptual framework regarding the reform's expected effect, based on Zweimüller (2018). The framework illustrates that the overall effect of the policy change can be decomposed into a *Behavioral Effect* due to changed individual-level search effort and a *Tightness Effect* due to changing labor market conditions:[3]

$$\underbrace{e_{post}f(\theta_{post}) - e_{pre}f(\theta_{pre})}_{\text{Overall Effect}} = \underbrace{[e_{post} - e_{pre}]f(\theta_{pre})}_{\text{Behavioral Effect}} + \underbrace{e_{post}[f(\theta_{post}) - f(\theta_{pre})]}_{\text{Tightness Effect}} \quad (3.2)$$

In this expression, e denotes a job seeker's effort and $f(\theta)$ its return, which depends on labor market tightness θ. As θ describes the ratio of vacancies over total search effort in the labor market, it is also affected by the policy change.

The behavioral effect denotes the difference in job finding rates that is solely attributable to the change in search effort at the individual level, keeping labor market conditions constant. The tightness effect, in turn, denotes the difference in the (potential) job finding rates of job seekers in post-reform versus pre-reform labor markets, keeping individual search effort constant. That is, the tightness effect denotes the difference in the job finding rate that is solely attributable to the change in labor market conditions.

Based on the simple framework, we now investigate heterogeneity in the reform effect along three dimensions: (1) individual-level effort change; (2) pre-reform labor market tightness; and (3) competition from untreated job seekers.

Individual-level effort change. We first analyze how the treatment effect varies with respect to the expected change in individual search effort. As a proxy for the expected effort change, we use predicted application requirement levels. Intuitively, if job seeker characteristics predict a high level of job search requirements in the absence of a policy change, then these job seekers would have been more restricted in their job search without the policy change. We thus regress the average number of application requirements per month on a vector of job seeker characteristics (same as x_i in equation (3.1)), using only the pre-reform data of job seekers in Bern. Using the estimates from this regression, we predict the individual average number of monthly application requirements that a job seeker would have gotten had she registered in Bern before the policy change. In Appendix 3.D, we show how predicted and actual requirements correlate with several control variables and with each other.
Based on the predicted individual average number of monthly application requirements, we conduct a median split. Figure 3.4 plots the *actual* individual average number of monthly application requirements for job seekers for whom we predicted a high level of application requirements and for those for whom we predicted a low level of application requirements. The drop in average required applications is larger for job seekers for whom we predict a high level of application requirements without

3. Note that we use this decomposition only to describe the different channels affecting the reform's effect. The variation in our data does not allow to separately estimate the behavioral effect and the tightness effect.

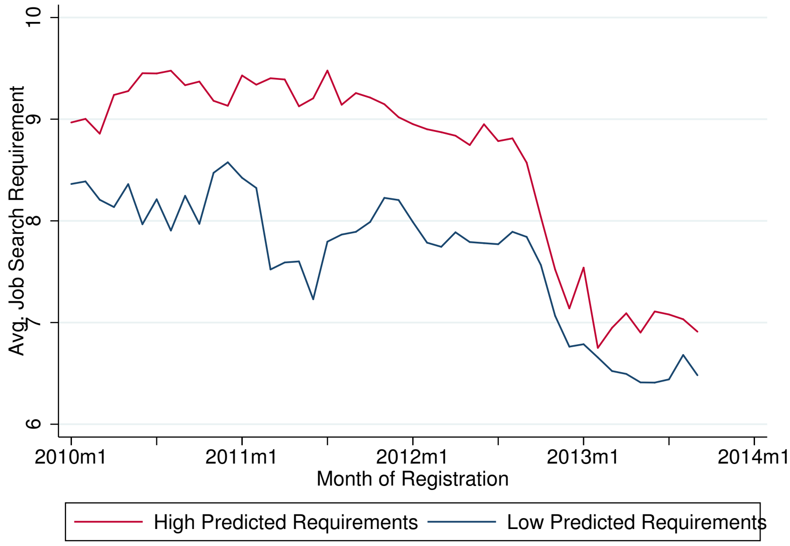

Figure 3.4. Application Requirements for High- vs Low-predicted Requirement Spells by Month of Registration into Unemployment

This figure displays how actual application requirements developed in Bern and shows the results from a median split based on the number of predicted monthly application requirements. We plot the average number of required applications per spell per month in Bern separately for UE-spell for which we predicted a high level of application requirements (red) and for those for which we predicted a low level of application requirements (blue), based on the pre-*BernTop!* data.

the policy change. That is, the job search effort for this group experiences a larger drop due to the reform.

Column (1) of Table 3.2 reports the results from a pooled Difference-in-Differences regression for individuals with high versus low predicted requirement levels (median split). The results show that the reform's effects are indeed about twice as strong for job seekers with a high versus low predicted pre-reform job search requirement level, supporting the idea that the size of the individual-level effort change is decisive for the reform's effect.

Pre-reform labor market tightness. As a second insight from Section 3.A, we expect the overall reform effect to increase with pre-reform local labor market tightness. Intuitively, effort changes matter more if there are more vacancies per job seeker available. To assess the relevance of this conjecture, we interact the Difference-in-Differences term with a median split on the pre-reform labor market tightness at the municipality level. Despite a lack of statistical significance, Column (1) of Table 3.2 shows suggestive evidence that the reform had a stronger effect in municipalities with a high tightness, i.e., where the ratio of vacancies over unemployed was more favorable from the job seekers' perspective.

Table 3.2. Effect on Unemployment Duration: Heterogeneity

	UE Duration		
	(1)	(2)	(3)
DiD	9.785**	11.005**	12.567***
	(4.016)	(4.334)	(3.473)
DiD x (High treatment intensity)	9.462*		
	(5.231)		
DiD x (High LM tighness)		6.524	
		(4.799)	
DiD x (Close to border)			8.467**
			(4.043)
Controls	Yes	Yes	Yes
Outcome Mean	266.27	266.27	266.27
R^2	0.114	0.116	0.114
N	348'448	348'431	348'431

The table reports estimates from a pooled version of the Difference-in-Differences framework defined by equation (3.1), excluding job seekers who registered in the year 2012. The outcome is the average unemployment duration in days (top-coded at 520 days). In column (1), "High treatment intensity" equals one if a job seeker's predicted requirement without *BernTop!* is at or above the median. In column (2), "High LM tightness" equals one if the vacancy-to-job seeker ratio in a job seeker's local labor market region is at or above the median. In column (3), "Close to border" equals one if the job seeker's municipality has a travel time of 15 minutes or less to the cantonal border. Standard errors are clustered at the PES level (N=150). *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Change in labor market tightness. The tightness effect sketched out in Section 3.A arises because job seekers in Bern collectively decrease their job search effort, positively affecting the probability of the individual job seeker to find a job and thus counteracting the behavioral effect. However, the smaller the share of job seekers in a labor market who decrease search effort, the smaller should the impact of this channel be.

As all job seekers who live in Bern are treated by the reform, an interesting source of variation in the share of treated job seekers stems from job seekers from other cantons searching for a job in Bern. We expect the tightness effect to matter more in labor markets with less competition from job seekers outside of Bern, who do not decrease their search. Intuitively, consider a closed labor market with no competition from job seekers from outside of the canton. Then a policy change as *BernTop!* would affect all job seekers in this closed labor market, who would collectively reduce their job search effort following the introduction of the policy change.

In contrast, consider Bern to be an open labor market where job seekers living in Bern compete with job seekers from neighboring cantons due to the possibility of commuting.[4] In such an open labor market, we expect the tightness effect to be smaller, as not all job seekers are treated by the reform.

While we cannot estimate the reform's effect under the scenario of a closed labor market, we can make use of the variation in the local competition from job seekers outside Bern. The simplest source of variation is the commuting time from a job seeker's home to the border of Bern and to one of the adjacent cantons. The calculation of commuting times is explained in Appendix 3.F. In the center of Bern, where commuting times to other cantons are longer, the share of job seekers from other cantons who are competing for the same jobs is lower. Indeed, as Figure 3.B.4 in the appendix shows, travel time and the share of commuters from the adjacent cantons are significantly negatively correlated.

To exploit this variation, we perform a split of the commuting time to the border of the canton. As before, we interact this split with the Difference-in-Differences coefficient. Column (3) of Table 3.2 shows that the reform had an about 30% stronger effect on job seekers in municipalities with a travel time of 15 minutes or less to the border. This result supports the idea that the average effect of policies that target search effort depend on the relative importance of tightness effects.

3.4.4 Externalities on Job Seekers in Adjacent Cantons

Given the finding that job seekers in Bern are more affected if they live closer to the border, it is natural to ask whether job seekers at the other side of the border were positively affected by the policy change. Intuitively, if job seekers in Bern decrease their effort, this could have positive externalities on the job finding prospects of job seekers outside Bern, who partially compete with job seekers from Bern. To test for potential externalities, we estimate the following regression on the population of job seekers outside Bern:

$$Y_i \;=\; \sum_{s=2009}^{2015} \gamma_s^{close}\, \mathbf{I}_{(y=s\,\&\,close=1)} + \pi_m + \lambda_{y\times c} + \sigma_t + X_i'\kappa + \varepsilon_i \tag{3.3}$$

The indicators $\mathbf{I}_{(y=s\,\&\,close=1)}$ equal one if a job seeker enters unemployment in year s and lives in a municipality with low commuting time to Bern.[5] The calculation of commuting times is explained in Appendix 3.F. π_m includes municipality fixed effects and $\lambda_{y\times c}$ are canton-year fixed effects. As in equation (3.1), σ_t are calendar

4. The share of workers living in Bern and commuting to a different canton is 10%. The share of commuters from other cantons working in Bern is 14%. Source: https://www.bve.be.ch/bve/de/index/mobilitaet/mobilitaet_verkehr/mobilitaet/grundlagen_mobilitaet/Pendlerstatistik.html, last visited: Jan 17, 2021

5. Figure 3.B.5 shows that commuting time and the share of job seekers in a municipality, who commute to the canton of Bern, are significantly negatively correlated.

month fixed effects and X_i are job seeker covariates. In this specification, we cluster standard errors at the level of the municipality, which defines the travel time to Bern.

The framework compares how the outcomes of job seekers close to the border evolve over time, relative to those of job seekers in the same canton, but further away from the border. This also means that cantons that are non-adjacent to Bern do not serve as a control group, but only identify the effect of covariates and aggregate time factors.

Figure 3.5 plots the estimates of γ_s^{close}, where "close" is defined by municipalities with a travel time of 30 minutes or less to the border of Bern and has the same language as the bordering municipality in Bern.[6] Although imprecisely estimated, the coefficients point towards a slight decrease in the average unemployment duration of job seekers in closeby municipalities. This is confirmed by the estimates from a pooled version of equation (3.3), as reported in Table 3.3.

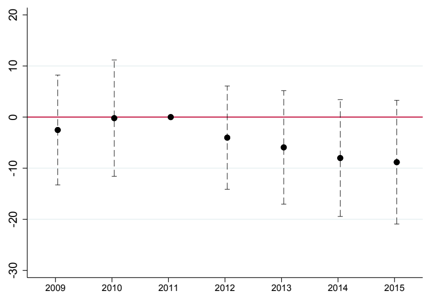

Figure 3.5. Effect on UE Duration of Job Seekers in Adjacent Municipalities ($\leq$ 30 minutes travel time to Bern)

Estimated coefficients of a dynamic Diff-in-Diff framework (equation (3.3)). The outcome is the duration of unemployment (top-coded at 520 days). Spells registered in Bern are excluded. The treatment group includes spells located in a municipality that is no more than 30 minutes (by car) away from the border of the canton of Bern and coincides with the language spoken in the corresponding part of Bern. The corresponding municipalities are displayed in Figure 3.B.2b. The dashed lines denote 95% confidence intervals. Standard errors are clustered at the municipality level. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

3.4.5 Effect on Average Earnings

So far, the results have shown that the increased autonomy in job search prolongued the average duration of unemployment spells. For a comprehensive evaluation of the reform and its monetary trade-offs, it is key to also understand how job match

6. Most often, this requires a municipality to be German speaking. An exception are muncipalities in the Cantons Jura and Neuchatel, which border the French-speaking part of Bern. In Appendix 3.F, we describe in detail how the commuting times from each municipality to the border of Bern were calculated.

Table 3.3. Effect on Unemployment Duration: Adjacent Cantons

	≥20 Mins	≥30 Mins	≥40 Mins
	(1)	(2)	(3)
DiD	0.727	-7.076*	-6.727**
	(4.332)	(3.626)	(3.255)
Controls	Yes	Yes	Yes
Outcome Mean	268.21	268.21	268.21
R^2	0.085	0.085	0.085
N	317’120	317’120	317’120

Notes: This table shows the estimated coefficients based on a pooled Diff-in-Diff regression. For this specification, we excluded observations from the canton of Bern and instead only used unemployment spells from the remaining cantons. As before, the outcome for this specification is unemployment duration (top-coded at 520 days). To produce the displayed coefficients, we generate a DiD-term based on a dummy variable that takes on the value 1 if an unemployment spell is located in a municipality that is no more than 20, 30 or 40 minutes (by car) away from the border of the canton of Bern and coincides with the language spoken in the corresponding part of Bern. The corresponding municipalities are displayed in Figure 3.B.2b. Standard errors are clustered at the municipality-level.
*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

quality was affected: did job seekers use their autonomy to successfully direct search towards longer-lasting, higher-paying jobs?

To shed light on this question, we study how post-unemployment earnings were affected by the reform. The effect of increased search autonomy on re-employment earnings is not easy to predict, even if the effect on the duration of unemployment is known. In a model with directed search (e.g. Nekoei and Weber, 2017), an increased duration of unemployment may reflect job seekers becoming more selective about the jobs they apply to. This would on average result in higher post-unemployment earnings. On the other hand, a longer unemployment duration has been shown to decrease the wage offers. With both effects working in opposite directions, the net effect is an empirical question.

We estimate the reform’s effect on post-unemployment earnings by using the same dynamic difference-in-differences model as outlined in equation (3.1). Instead of unemployment duration, we now use log re-employment earnings from social security records as the outcome. Figure 3.6 presents the results. We estimate an increase in post-unemployment earnings between two and four log points on average –pointing towards the dominance of the first mechanism. It appears that search autonomy can yield the benefit of allowing individuals to find better job matches. Unfortunately, we have a limited observation window after the reform was implemented and post-unemployment effects started to show. We thus cannot inform about the persistence of these earnings effects. Extending the time window will be an important part of ongoing and future work.

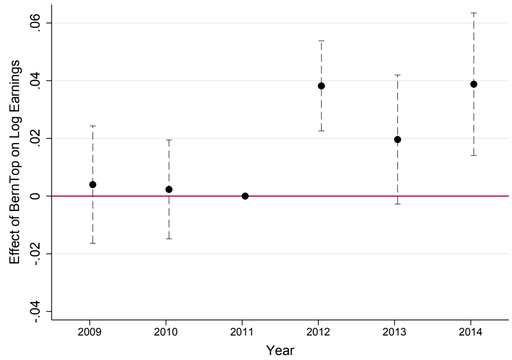

Figure 3.6. Estimated Dynamic Effect on log Post-Unemployment Earnings

This figure displays the estimated coefficients based on the dynamic Diff-in-Diff regression (3.1). The outcome is the log of post-unemployment monthly employment earnings. The dashed lines denote 95% confidence intervals. Standard errors are clustered at the PES level.

Table 3.4 presents the results from pooled Difference-in-Differences regressions corresponding to the results in Figure 3.6. The pooled regression delivers an estimate of the earning effect of 2.4 log points. Additionally, we estimate a decrease in the probability to incur an earnings loss compared to the pre-unemployment earnings of 2.6 percentage points. Finally, we take a look at post-unemployment job stability: We estimate the effect of the reform on the probability to register as unemployed again after less than 6 months. We find suggestive evidence that post-unemployment job stability increased following the reform, though the estimate is statistically insignificant. Note, however, that the post-unemployment analysis of job stability is subject to a limited observation window, that extends until late 2015. Thus, post-unemployment job stability can only be assessed for job seekers entering unemployment shortly after the reform, which significantly decreases the size of the treatment group. Table 3.E.2 in the appendix shows that the estimated effect on post-unemployment earnings (and its statistical significance) is robust with respect to bootstrapped standard errors, clustering of the standard errors at the canton level, including spells that entered unemployment in 2012, including canton- or RAV-fixed effects and restricting the control group to zipcode areas with more than 40 minutes commuting time to the canton of Bern.

Table 3.G.2 presents results for several job seeker sub groups. We find that the positive earnings effect is mostly driven by female job seekers, job seekers in White Collar occupations and the Service Sector and job seekers of Swiss Nationality. In particular, the earnings effect appears to be especially strong for those groups of job seekers for whom we estimated a weak average effect on unemployment duration.

Table 3.4. Estimated Pooled Effect on Post-Unemployment Outcomes

	Log Earnings		P(Earnings Loss)	P(Empl. < 6 Mon.)
	(1)	(2)	(3)	(4)
DiD	0.027***	0.024***	-0.026**	-0.006
	(0.010)	(0.009)	(0.011)	(0.008)
Controls for Prev. Earnings	Yes	Yes	Yes	Yes
Controls	No	Yes	Yes	Yes
Outcome Mean	8.47	8.47	0.54	0.16
R^2	0.372	0.396	0.132	0.021
N	187’490	187’490	187’490	213’358

Notes: This table reports estimates from a pooled version of the Difference-in-Differences framework defined by equation (3.1), excluding job seekers who registered in the year 2012. Additionally, we restrict the sample to job seekers who entered unemployment before July 2014 and found a job in less than 520 days. The outcome in columns (1) and (2) is the log of post-unemployment monthly employment earnings. In column (3), the outcome is the probability to enter a job with lower monthly earnings than the average monthly earnings during the 12 months prior to unemployment. In column (4), the outcome is the probability to exit employment after less than 6 months. That is, in this column, we analyze the effect on job stability. All specifications control for previous earnings. Standard errors are clustered at the PES level (N=150).
*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

3.5 Monetary Trade-Offs

3.5.1 Fiscal Trade-Off

When estimating the average effect of a policy change on unemployment duration and post-unemployment earnings, a natural follow-up analysis concerns the fiscal trade-off inherent to the policy change. A meaningful analysis of the fiscal gains or losses encompasses multiple components we need to consider: The effect on unemployment, the effect on post-unemployment earnings, the effect on post-unemployment job stability and the persistence of all these effects. Since our observation window covers unemployment spells only up to late 2015, any analysis of post-unemployment job stability or the persistence of the effect of the policy change on earnings would have limited informative value. However, our estimates regarding unemployment duration and post-unemployment earnings allow for a back-of-the-envelope calculation and provide a starting point for a fiscal cost-benefit analysis. Note that for this calculation, we ignore any externalities in the remaining cantons in Switzerland and instead only focus on the canton of Bern. The fiscal costs the average job seeker generates from staying unemployed for 14.5 additional days are[7]

7. Note that we use the data on pre-unemployment earnings and replacement rate to retrieve the average daily wage and replacement rate in our sample.

$$\underbrace{14.5}_{\text{Effect on avg. UE duration}} \times \underbrace{165\ CHF}_{\text{avg. daily wage}} \times \underbrace{(0.75}_{\text{avg. benefit rate}} + \underbrace{0.2}_{\text{avg. income tax}} (1-0.75))$$
$$= 1'914\ CHF$$

That is, a job seeker staying unemployed 14.5 days longer (earning 165 CHF per day on average if she were not unemployed) results in forgone tax revenues of $14.5 \times 165 \times 0.2 = 495$ CHF. Through unemployment benefits the fiscal costs increase by $15 \times 165 \times 0.75 = 1856.25$ CHF, based on an average replacement rate of 0.75. On the other hand, job seekers in Switzerland pay taxes on unemployment benefits, resulting in tax revenues of $15 \times 165 \times 0.75 \times 0.2 = 371.25$ CHF.
In contrast, the fiscal gain from job seekers finding higher-paying jobs can be quantified using the effect on job stability, the effect on post-unemployment earnings and its persistence. We estimate that job seekers experience an increase in post-unemployment earnings of 2.4 log points:

$$\underbrace{0.024}_{\text{Effect on avg. wage}} \times \underbrace{165\ CHF}_{\text{avg. daily wage}} \times \underbrace{0.2}_{\text{avg. income tax}} = 0.792\ CHF$$

That is, the additional daily fiscal gains from income taxation through increased income amount to 0.79 CHF per day on average. As our data does not allow for a meaningful analysis regarding job stability or the long-term persistence of the wage effect, we cannot include them in this calculation. Instead, we assume that the estimated effects persist indefinitely and exclude the notion of job stability from our analysis. Instead of quantifying the effect of the reform on job stability, we calculate how long it would take the fiscal gains of the reform to break even with its fiscal costs. With average fiscal costs of 1'914 CHF per spell and average daily fiscal gains of 0.792 CHF, we calculate that the earnings effect would (in the absence of an effect on job stability) cancel out the fiscal costs after $1'914/0.792 = 2'417$ days (≈ 6.6 years).

3.5.2 Worker Trade-Off

We now complement the fiscal trade-off with the monetary trade-offs faced by individual workers. While workers also incur forgone income from longer unemployment, the benefits from finding a better paying job carry a relatively higher weight compared to the forgone income than in the fiscal trade-off. A job seeker's costs from remaining 14.5 more days in unemployment can again be calculated from the estimation results in Table 3.1:

$$165\ CHF \times [(1-0.2)(1-0.75)] \times 14.5 = 478.5\ CHF$$

That is, the average job seeker incurs forgone income of 165 $CHF \times (1-0.2) \times 14.5 = 1'914\,CHF$, but gets compensated with a 75 percent replacement rate: $1'914\,CHF \times (1-0.75) = 478.5\,CHF$. After leaving unemployment, job seekers experience income gains of 2.4 log points on average. This results in daily gains of:

$$165 \;\; CHF \times 0.024 \times (1 - \underbrace{0.2}_{\text{avg. income tax}}) \; = \; 3.168 \;\; CHF$$

As before, we can calculate how long it would take a job seekers income gains to cancel out the forgone income from staying unemployed for 14.5 more days on average. These are $478.5\,CHFs/3.168\,CHFs\;per\;day \approx 151\;days$. This indicates that the income gains amortize the forgone income after 5 months.

3.6 Conclusion

We analyze the effects of increasing the job search autonomy of unemployed workers. We exploit a policy change in the Swiss canton Bern, during which application requirements were drastically reduced and the use of official job referrals was almost completely abolished. We find that, on average, unemployment duration increased by about 8%. The effect is larger among job seekers whose effort is predicted to reduce more strongly due to the reform. Moreover, we find that the changes in local labor market tightness due to the fact that job seekers collectively decrease their effort mitigate the effect of individual effort changes.

Finally, we present evidence on the reform's effect on post-unemployment earnings. We estimate that post-unemployment earnings increased by 2 to 4% on average. Apparently, the decreased pressure through increased search autonomy allowed job seekers to find slightly better matching and better paying jobs.

In sum, our estimates indicate that increasing the autonomy and self-responsibility of job seekers carries the risk of prolonging unemployment spells, but these adverse effects can, from a fiscal point-of-view, be eventually compensated by an increase in post-unemployment earnings. Furthermore, results suggest that the success of policies which lead to collective changes in search effort depends on the local scope for tightness effects and search externalities. In the context of this study, the behavioral effect of a decrease in individual effort was apparently the dominating force. It is open whether this also holds in other labor markets with less favorable labor demand conditions.

Appendix 3.A Conceptual Framework

Setup. Let *post* denote the labor market after the reform and *pre* denote the labor market before the reform. We start by assuming the canton Bern to be a closed labor market, such that all job seekers in the *post* labor market are treated. Let V denote the number of vacancies and E the aggregate amount of search effort provided by job seekers in Bern. Thereby, $\theta = V/E \ \in \{\theta_{post}, \theta_{pre}\}$ describes labor market tightness in the *post* and *pre* labor market, respectively. The individual search effort of *post* and *pre* job seekers is defined by e_{post} and e_{pre}. Assuming that every job offer is accepted, the job finding probability $e \times f(\theta)$ depends on individual job search effort e and a function f that is increasing in labor market tightness θ, with $f''(\cdot) < 0$. If the reform induces job seekers to exert less (but still positive) search effort, we have $e_{post} < e_{pre}$.

The overall effect of the policy change can be decomposed into a *Behavioral Effect* due to changed individual-level search effort and a *Tightness Effect* due to changing labor market conditions.

$$\underbrace{e_{post}f(\theta_{post}) - e_{pre}f(\theta_{pre})}_{\text{Overall Effect}} = \underbrace{[e_{post} - e_{pre}]f(\theta_{pre})}_{\text{Behavioral Effect}} + \underbrace{e_{post}[f(\theta_{post}) - f(\theta_{pre})]}_{\text{Tightness Effect}} \quad (3.A.1)$$

The behavioral effect denotes the difference in job finding rates that is solely attributable to the change in search effort at the individual level, keeping labor market conditions constant. The tightness effect, in turn, denotes the difference in the (potential) job finding rates of job seekers in post-reform versus pre-reform labor markets, keeping individual search effort constant. That is, the tightness effect denotes the difference in the job finding rate that is solely attributable to the change in labor market conditions.

The relationship between the overall effect, the behavioral effect and the tightness effect sketched out in equation (3.A.1) is illustrated in Figure 3.A.1. We plot the job finding rate as a function of the job search effort e and the labor market tightness θ. In tighter labor markets (high θ), job seekers face less competition for the existing vacancies, which increases the individual job finding rate. Hence, the job finding rate increases in θ. Secondly, the job finding rate increases in the exerted search effort. Thus, when job seekers decrease search effort following the introduction of *BernTop!* ($e_{post} < e_{pre}$), the individual job finding rate decreases for every potential labor market tightness, as depicted by the blue graph staying below the red graph for all values of θ. Finally, we assume that the returns of labor market tightness to the individual job finding rate diminish, resulting in a concave function.

The (negative) behavioral effect of the policy reform denotes the difference in the job finding rate in the pre-reform labor market between treated and untreated job seekers, that is, the effect of decreasing job search following the introduction of

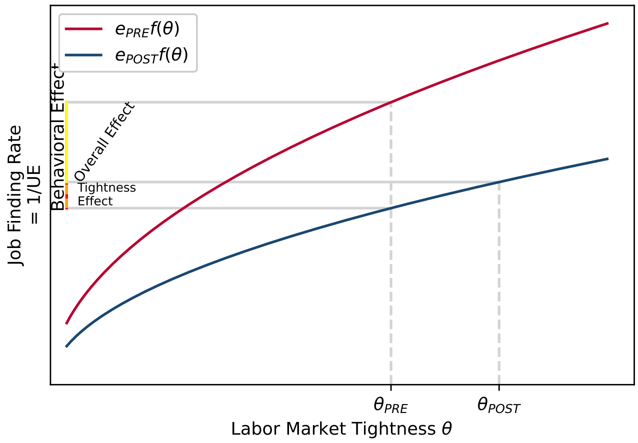

Figure 3.A.1. Decomposition of Reform Effect

We display the graphical decomposition of the overall effect of a policy change that decreased job search effort for all job seekers in a closed economy. The (negative) behavioral effect denotes the decrease in the job finding rate that is purely attributable to the decrease in job search effort from e_{PRE} to e_{POST}, evaluated at the pre-reform market tightness θ_{PRE}. The (positive) tightness effect denotes the increase in the job finding rate that is purely attributable to changing the labor market tightness from θ_{PRE} to θ_{POST}, evaluated at the post-reform search effort e_{POST}.

BernTop!. The tightness effect denotes the difference between the job finding rate of a post-reform job seeker in a post-reform and a pre-reform labor market, that is, the effect of changing labor market conditions, while keeping individual search effort constant. The overall effect is the sum of the two effects, which work in opposite directions.

Insights for the Empirical Analysis. This simple decomposition yields insights regarding the expected effect of the policy change under different labor market conditions.

First, it is straight-forward to expect that the overall effect on job finding will be more pronounced for job seekers who reduced their search effort more than others due to the reform, as we expect a stronger behavioral effect for these job seekers.

Second, we expect the policy change to have a weaker effect in slack labor markets, where many job seekers compete for few vacancies, than in tight labor markets with an abundance of vacancies. Intuitively, reducing search effort in a slack labor market with few vacancies per job seeker does not decrease the job finding rate as much as in a tight labor market. This scenario is similar to a game of musical chairs, in which many people try to sit on far too few chairs. If one person exerts less effort to sit on one of the free chairs, the probability to get the free chair, which was already low to begin with, did not decrease much. This scenario is illustrated in Figure 3.A.2, panel (a). The behavioral effect of reducing job search effort from e_{pre} to e_{post} on the job finding rate is smaller in slack labor markets than in tight labor markets.

The derivative of the expected overall effect with respect to pre-reform labor market tightness θ_{pre} is derived in Section 3.A.1.

Third, the framework implies that the overall effect decreases with the change in local labor market tightness due to the reform.[8] For illustration, consider a labor market in which only some job seekers are treated by the reform. The lower the share of untreated job seekers, the higher the share of job seekers who decrease search effort. Hence, we expect the tightness effect to play a relatively larger role in labor markets with a high share of treated job seekers. In these markets, we expect the average effect of the policy change on treated job seekers to be lower. This is illustrated in Figure 3.A.2, panel (b). If the difference in the labor market tightness between treated and untreated labor markets is large, the (positive) tightness effect negates the (negative) behavioral effect.

8. This result is derived in Appendix 3.A.1 under the assumption that the change in labor market tightness $\theta_{post} - \theta_{pre}$ consists **only** of variation in pre- **or** post-reform labor market tightness, keeping post- and pre-reform labor market tightness constant, respectively.

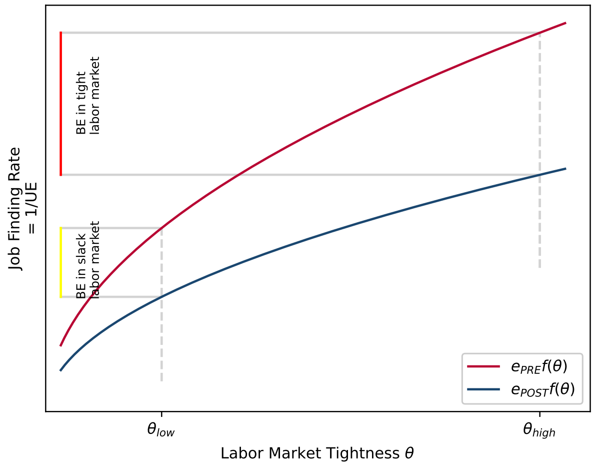

(a) Micro Effect in Tight vs. Slack Labor Markets

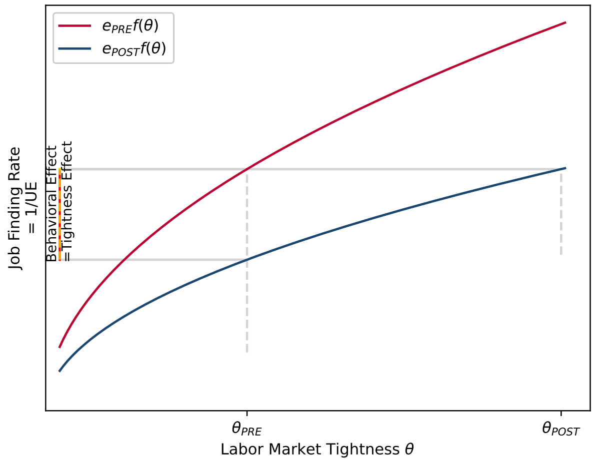

(b) The Positive Externality Offsets the Negative Micro Effect

Figure 3.A.2. Effect of Labor Market Conditions

This figure show the expected effect of the reform in different labor market conditions. Panel (a) displays the expected behavioral effect in tight vs. slack labor markets, that is, in competitive vs. less competitive labor markets (from a job seeker's perspective). Panel (b) shows how the overall estimated effect can change, depending on the change in labor market conditions that comes along with the policy change.

3.A.1 Derivations

We assume that $f' > 0, f'' < 0$.

$$\begin{aligned}
\underbrace{OE}_{\text{Overall Effect}} &= [e_{post} - e_{pre}]f(\theta_{pre}) + e_{post}[f(\theta_{post}) - f(\theta_{pre})] \\
\Rightarrow \frac{\partial OE}{\partial \theta_{pre}} &= [e_{post} - e_{pre}]f'(\theta_{pre}) + e_{post}\left[f'(\theta_{post})\frac{\partial \theta_{post}}{\partial \theta_{pre}} - f'(\theta_{pre})\right]
\end{aligned}$$

We will assume that $\frac{\partial \theta_{post}}{\partial \theta_{pre}} = 1$ and thus

$$\frac{\partial OE}{\partial \theta_{pre}} = \underbrace{\underbrace{[e_{post} - e_{pre}]}_{<0}\underbrace{f'(\theta_{pre})}_{>0} + e_{post}\left[\underbrace{f'(\theta_{post}) - f'(\theta_{pre})}_{<0}\right]}_{<0}$$

Note that the TE is negative, so a negative derivative indicates in increase in magnitude of the TE. But also important:

$$\frac{\partial OE}{\partial \Delta\theta} = \underbrace{[e_{post} - e_{pre}]}_{<0}\underbrace{f'(\theta_{pre})}_{>0}\frac{\partial \theta_{pre}}{\partial \Delta\theta} + e_{post}\left[\underbrace{f'(\theta_{post})}_{>0}\frac{\partial \theta_{post}}{\partial \Delta\theta} - \underbrace{f'(\theta_{pre})}_{>0}\frac{\partial \theta_{pre}}{\partial \Delta\theta}\right] \tag{3.A.2}$$

With $\theta_{post} - \theta_{pre} > 0$, we will/should assume that $\frac{\partial \theta_{post}}{\partial \Delta\theta} \geq 0$ and $\frac{\partial \theta_{pre}}{\partial \Delta\theta} \leq 0$ and thus

$$\frac{\partial OE}{\partial \Delta\theta} = \underbrace{\underbrace{[e_{post} - e_{pre}]}_{<0}\underbrace{f'(\theta_{pre})}_{>0}\underbrace{\frac{\partial \theta_{pre}}{\partial \Delta\theta}}_{\geq 0}}_{\leq 0} + \underbrace{e_{post}\left[\underbrace{f'(\theta_{post})}_{>0}\underbrace{\frac{\partial \theta_{post}}{\partial \Delta\theta}}_{\geq 0} - \underbrace{f'(\theta_{pre})}_{>0}\underbrace{\frac{\partial \theta_{pre}}{\partial \Delta\theta}}_{\leq 0}\right]}_{\geq 0} \tag{3.A.3}$$

Finally, we can simplify the expression by assuming $\frac{\partial \theta_{pre}}{\partial \Delta\theta} = 0$ or $\frac{\partial \theta_{post}}{\partial \Delta\theta} = 0$, respectively, such that either

$$\frac{\partial OE}{\partial \Delta\theta} = \underbrace{e_{post}\underbrace{f'(\theta_{post})}_{>0}\underbrace{\frac{\partial \theta_{post}}{\partial \Delta\theta}}_{\geq 0}}_{\geq 0} \tag{3.A.4}$$

or

$$\frac{\partial OE}{\partial \Delta\theta} = \underbrace{\underbrace{[e_{post} - e_{pre}]}_{<0}\underbrace{f'(\theta_{pre})}_{>0}\underbrace{\frac{\partial \theta_{pre}}{\partial \Delta\theta}}_{\leq 0}}_{\geq 0} + \underbrace{e_{post}\left[-\underbrace{f'(\theta_{pre})}_{>0}\underbrace{\frac{\partial \theta_{pre}}{\partial \Delta\theta}}_{\leq 0}\right]}_{\geq 0} \tag{3.A.5}$$

Appendix 3.B Additional Figures

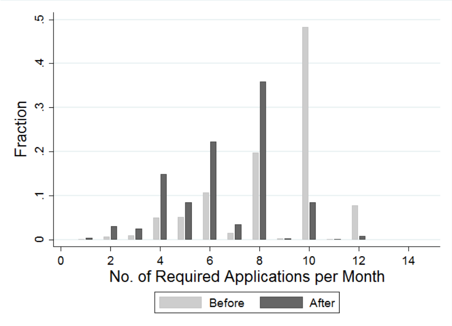

(a) Distribution of Required Search Effort in Bern

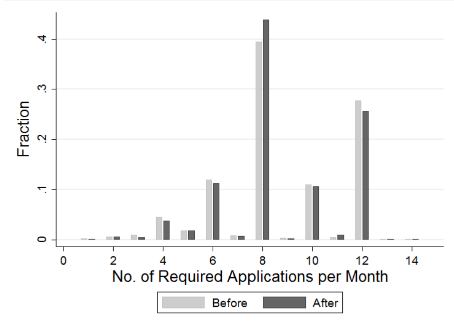

(b) Distribution of Required Search Effort in the Remaining Cantons

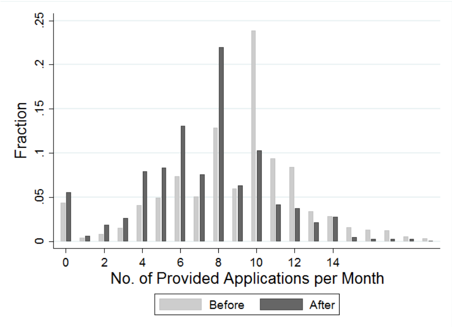

(c) Distribution of Reported Search Effort in Bern

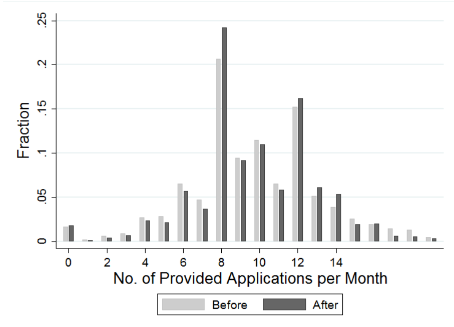

(d) Distribution of Reported Search Effort in the Remaining Cantons

Figure 3.B.1. Distribution of Search Effort

This figure shows the distribution of required and reported search effort (applications) in Bern and the other cantons with data on required and provided effort (Fribourg, Solothurn and Tessin), before and after the reform was implemented.

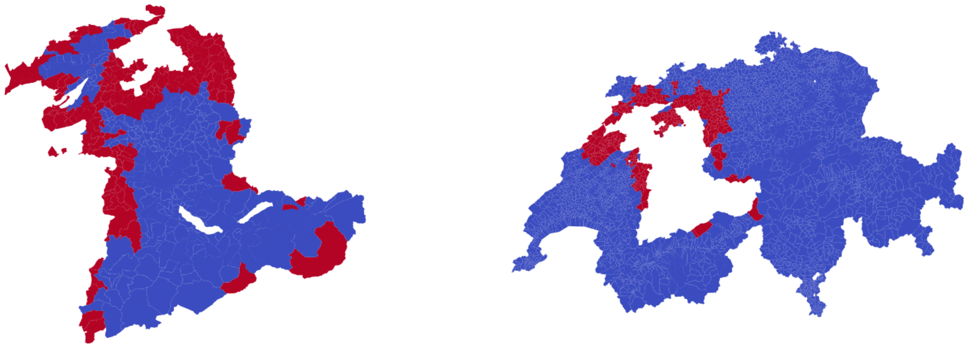

(a) Closeness to Border in Bern **(b)** Closeness to Border in the Adjacent Cantons

Figure 3.B.2. Closeness to Border

This figure shows the zipcode areas in our sample which were used for the distance analyses presented in Section 3.4.4. In Figure 3.B.2a, we show the canton of Bern and mark the zipcode areas close to the border of Bern. We split the zipcode areas by travel time to the next canton, according to Google Maps. We define a zipcode area as "close" to the border if it takes less than 15 minutes to reach the border by car. In Figure 3.B.2b we mark the zipcode areas close to the canton of Bern. We define a zipcode area as "close" to the canton of Bern, if it takes less than 30 minutes to reach the canton of Bern by car and if the language region in the zipcode area coincides with the language region in the corresponding part in Bern. Using this additional criterion, we make sure that these commuters indeed search in the same labor markets as the commuters in Bern and are not restricted by language barriers. Our split defines the southern regions in Bern as well as the regions south of Bern as "not close". This is due to the "Bernese Alps", a mountain range in the south of Bern, which only leaves a handful of passes between the southern canton of Valais and Bern and thus makes commuting between both cantons very time consuming.

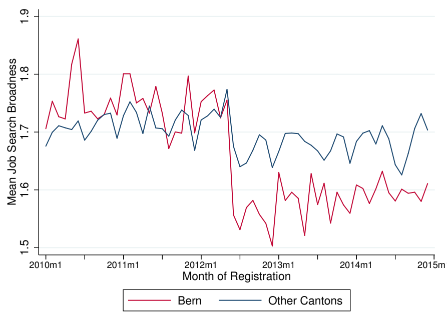

Figure 3.B.3. Broadness of Job Search

This figure displays the average broadness in the classified occupations job seekers declared to search in. Occupational broadness is defined as the variation in the first two digits of the BN-2000 codes of the occupations a job seeker states to search in. *BernTop!* also prescribed caseworkers not to push job seekers towards searching in occupations too far outside their experience and competence. During the first meeting with the caseworkers, job seekers fill in a form stating in which occupations they are searching for a job. These occupations are then classified according to the BN-2000 nomenclature, which classifies occupations using a 5-digit code. Each of the first three digits further specifies an occupation in a lexicographic fashion, while the last two digits pin it down. For example, the occupation of a mason is classified by the code "41101" as follows: "4" denotes "Occupations in Construction & Mining", "4.1" denotes "Occupations in Construction", "4.1.1" denotes "Occupations in Core Construction" and "4.1.1.01" denotes "Masons". Hence, the more similar two occupations are, the more of the first digits of their BN-2000 codes coincide.

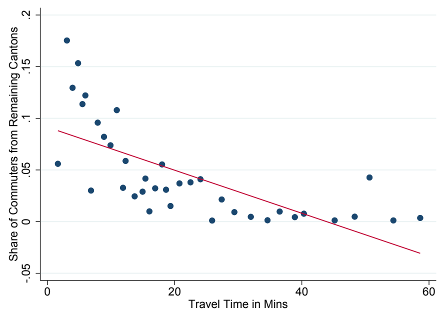

Figure 3.B.4. Travel Time and Share of Outside Commuters

This figure shows a scatterplot of the commuting time to the border of Bern in minutes on the x-axis to the share of commuters from other cantons for each municipality in Bern. For clarity reasons, we generated bins of travel times. For the travel times, we used the same weighted distances based on the zipcode travel times we used in the externality Section 3.4.4. The correlation coefficient between both variables is -0.258 (p-value=0.000).

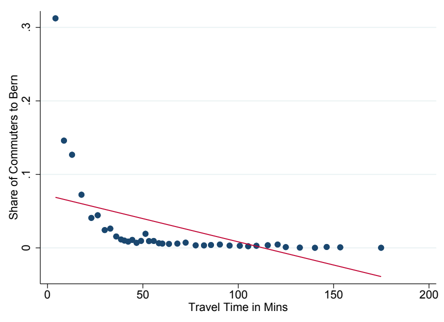

Figure 3.B.5. Travel Time and Share of Commuters to Bern

This figure shows a scatterplot of the commuting time for each municipality in the remaining cantons except Bern to the border of Bern in minutes on the x-axis to the share of commuters from that municipality to Bern. For clarity reasons, we generated bins of travel times. For the travel times, we used the same weighted distances based on the zipcode travel times we used in the externality Section 3.4.4. The correlation coefficient between both variables is -0.426 (p-value=0.000).

Appendix 3.C Summary Statistics

Table 3.C.1. Summary Statistics of Key Covariates

	Mean	Min	Max	Observations
Age	40.965	28	60	417’780
Age Squared	1757.216	784	3600	417’780
No other Person Affected[a]	0.650	0	1	417’780
One Person Affected[a]	0.158	0	1	417’780
One to Three Persons Affected[a]	0.175	0	1	417’780
More than Three Persons Affected[a]	0.017	0	1	417’780
German Language Proficiency Score	2.217	1	4	417’780
French Language Proficiency Score	2.666	1	4	417’780
Sex[b]	0.353	0	1	417’780
Low Education[c]	0.285	0	1	417’780
Medium Education[c]	0.392	0	1	417’780
High Education[c]	0.323	0	1	417’780
Potential Benefits Duration 400 Days	0.163	0	1	417’780
Married	0.492	0	1	417’744
Single, Widowed or Divorced	0.508	0	1	417’780
Non-Swiss	0.486	0	1	417’780
Non-Permanent Resident	0.207	0	1	417’780
Employment in Last 24 Months prior to UE[d]	21.517	2	24	417’780
Employment in Last 30 Months prior to UE[d]	26.415	2	30	417’780
Log Average Wage over the last 24 Months	8.500	5	13	417’780
Log Average Wage over the last 30 Months	8.494	5	13	417’780
German language region	0.624	0	1	417’780
French language region	0.315	0	1	417’780
Italian language region	0.058	0	1	417’780
Rhaeto-Romanic language region	0.004	0	1	417’780
Occupation in Production	0.151	0	1	417’780
Occupation in Engineering	0.068	0	1	417’780
Occupation in Construction	0.132	0	1	417’780
Occupation in Sales	0.088	0	1	417’780
Occupation in Gastronomy/Tourism	0.212	0	1	417’780
Occupation in White Collar or Health	0.319	0	1	417’780

Notes: This table presents summary statistics for the covariates used in all previous specifications.
[a]: Number of persons affected = Household size - 1
[b]: 0: Male, 1: Female
[c]: Low Education: Max. Compulsory High School, Medium Education: Max. Vocational Training, High Education: Max. College
[d]: Number of Months of Employment During the Last X Months prior to Unemployment

Appendix 3.D Correlation Structures

Table 3.D.1. Correlation Structure Between Key Characteristics

	Actual Monthly Requirements	Predicted Monthly Requirements
Actual Monthly Requirements	1.000	
Predicted Monthly Requirements	0.396***	1.000
Age	-0.065***	-0.226***
Native Language: German	-0.214***	-0.421***
Native Language: French	-0.033***	0.020***
Nr. of Affected Persons	0.003	0.017***
Sex[a]	0.036***	0.093***
Low Education[b]	0.102***	0.261***
High Education[b]	-0.136***	-0.432***
Potential Benefit Duration[c]	-0.059***	-0.178***
Single	-0.028***	-0.005
Married	0.037***	0.036***
Widowed	0.006	0.007
Divorced	-0.015***	-0.047***
Non-Swiss	0.133***	0.308***
Employment in 6 Months before UE[d]	0.019***	0.054***
Employment in 12 Months before UE[d]	-0.062***	-0.139***
Employment in 24 Months before UE[d]	-0.086***	-0.214***

Notes: This table shows unconditional correlations between job seeker/unemployment spell characteristics and the predicted and actual number of application requirements per month set by the caseworkers. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

[a]: 0: Male, 1: Female

[b]: Low Education: Max. Compulsory High School, Medium Education: Max. Vocational Training, High Education: Max. College

[c]: In Days

[d]: In Months

3.D.1 t-test by Categorical Variables

Table 3.D.2. Predicted and Actual Requirements

	Predicted Requirements			Actual Requirements		
	0	1	p-value	0	1	p-value
Occupation[a]	8.941	8.241	0.000	8.393	8.018	0.000
Married[b]	8.568	8.616	0.000	8.182	8.303	0.000
Education[c]	9.260	7.574	0.000	8.620	7.785	0.000
Over 40 Years	8.838	8.333	0.000	8.405	8.050	0.000
$\leq$ 1 Person Affected	8.595	8.589	0.026	8.213	8.261	0.045
Native Speaker	8.872	8.400	0.000	8.228	8.249	0.379
Nationality[d]	8.311	8.888	0.000	7.920	8.588	0.000
Non-Permanent Resident	8.503	8.931	0.000	8.141	8.749	0.000

Notes: This table shows the average predicted- and average actual application requirements per month set by the caseworkers for job seekers with different characteristics. The p-values are based on t-tests on the equality of means.
[a]: 0: Blue Collar/Low Paid Service Sector, 1: White Collar
[b] 0: Single/Widowed/Divorced 1: Married
[c] 0: “High” education (Max. degree: College) 1: “Low” education (Max. degree: Compulsory high school)
[d] 0: Swiss, 1: Non-Swiss

Appendix 3.E Robustness Checks

Table 3.E.1. Robustness Checks for the Main Difference-in-Differences Results

	Pooled DiD	With Bootstrapped SEs	Clustering on Canton Level	With 2012	With Canton FEs	With RAV FEs	>40 Mins Traveltime
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
DiD	14.472***	14.472***	14.472***	14.601***	12.489***	10.661***	13.752***
	(3.728)	(2.270)	(3.498)	(3.358)	(3.675)	(3.352)	(3.828)
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Outcome Mean	266.27	266.27	266.27	265.98	266.27	266.27	263.67
R^2	0.114	0.114	0.114	0.113	0.142	0.153	0.115
N	348’448	348’448	348’448	417’780	348’448	348’448	300’508

Notes: This table shows the results from several pooled Difference-in-Differences regressions. The basis for the analyses are the regressions shown in Table 3.1. We treat job seekers from Bern, who entered unemployment in 2012 or later in Bern, as treated. Job seekers entering unemployment in 2012 or later in the remaining cantons are used as the control group. We use the same set of covariates used in equation (3.1). In column (1), we show the estimated DiD-coefficients from the basic pooled Diff-in-Diff analysis in Table 3.1. In column (2), we use bootstrapped standard errors (50 bootstrap replications). In column (3), we cluster all standard errors at the canton level. In column (4), we include observations from the year 2012, which we previously left out. In columns (5) and (6), we add canton- and RAV fixed effects (in addition to controlling for the canton of Bern in the DiD-framework). In column (7), we only include observations from Bern and from municipalities with a commuting time longer than 40 minutes.

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

Table 3.E.2. Robustness Checks for the Difference-in-Differences Earnings Results

	Pooled DiD	With Bootstrapped SEs	Clustering on Canton Level	With 2012	With Canton FEs	With RAV FEs	>40 Mins Traveltime
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
DiD	0.024***	0.024***	0.024***	0.023***	0.025***	0.024***	0.024***
	(0.009)	(0.007)	(0.003)	(0.006)	(0.009)	(0.009)	(0.009)
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Outcome Mean	8.47	8.47	8.47	8.47	8.47	8.47	8.47
R^2	0.396	0.396	0.396	0.400	0.396	0.398	0.401
N	187’490	187’490	187’490	234’864	187’490	187’490	162’359

Notes: This table shows the results from several pooled Difference-in-Differences regressions. The basis for the analyses are the pooled DiD regressions on post-unemployment wages shown in Table 3.4. We treat job seekers from Bern, who entered unemployment in 2012 or later in Bern, as treated. Job seekers entering unemployment in 2012 or later in the remaining cantons are used as the control group. We use the same set of covariates used in equation (3.1). In column (1), we show the estimated DiD-coefficients from the basic pooled Diff-in-Diff analysis in Table 3.1. In column (2), we use bootstrapped standard errors (50 bootstrap replications). In column (3), we cluster all standard errors at the canton level. In column (4), we include observations from the year 2012, which we previously left out. In columns (5) and (6), we add canton- and RAV fixed effects (in addition to controlling for the canton of Bern in the DiD-framework). In column (7), we only include observations from Bern and from municipalities with a commuting time longer than 40 minutes.
*** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$

Appendix 3.F Distance Calculation

In this section, we describe in detail how the travel distances and travel times used for the analysis in Section 3.4.4 were calculated. We first calculate the travel distances and -times based on zipcode level and then use a weighted average of zipcode travel distances and -times for each municipality, at which the geolocation information for each unemployment spell is defined. The corresponding weights to assign zipcodes to municipalities are defined as the percentage of buildings in a municipality belonging to the corresponding zipcode area. The weights are provided by the Swiss Federal Office of Statistics.[9]

To calculate the travel distances and -times, we use a two-step-procedure: In the first stage, for each zipcode in Switzerland, we use the latitude/longitude-geolocation data generated by the `Nominatim()`-function contained in the `pgeocode` Python library to calculate the center point (or point of the largest village, respectively) for each zipcode area. Using the latitude/longitude-geolocation data and `geojson`-files containing geodata of the Swiss cantons, we calculate the shortest (air-)distance from a zipcode area to the border of Bern. Using this procedure, we generate the start- and end points to calculate travel distances and -times in a second step. Figure 3.F.1 illustrates how the closest point on the border of Bern (in this example for the zipcode area 3237 (Brüttelen)). Note that in the first stage of calculating travel distances and -times, we do not yet account for geographical special cases like mountains or road difficulties. However, executing the first stage with traveling distances and -times would require to calculate these for every point on the border of Bern, which is not feasible.

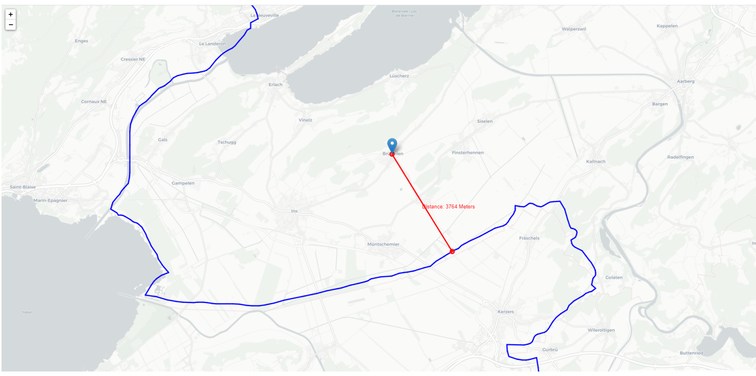

Figure 3.F.1. Calculation of Closest Point on Border

This figure shows an example of how the border point closest to a zipcode was calculated. The point delivered by `pgeocode` is the starting point for the calculation. We then calculate the air distance to every point on the border of Bern (Blue Line) and continue with the point for which this distance is the smallest.

9. https://www.bfs.admin.ch/bfs/de/home/statistiken/kataloge-datenbanken.assetdetail.7226419.html, last visited: Jan 10, 2022

In the second stage, we use the assignment of each zipcode centerpoint in Switzerland (both inside and outside of Bern) to the closest point on the Bern border to calculate the travel distance and travel time (by car) between both points using the Google Maps Distance Matrix API. This service by Google allows for large travel calculation queries for a series of locations. Using this approach, we can account for the special geography of Switzerland, with many mountain ranges and lakes that create a discrepancy between air distance and commuting distance. In particular, this approach reveals that, for example, the canton of Valais, despite sharing a large portion of the border, does not allow for quick commuting between both cantons. This is due to the "Bernese Alps", a mountain range in the south of Bern, which only leaves a handful of mountain passes between both cantons and thus makes commuting between them very time consuming. This mountain range is also the reason why, for some zipcodes, the travel distances and -times cannot be calculated using an API, because the closest point on the Bern border lies on a point in the mountain range, which cannot be reached by car. For these special cases, we calculate the travel distances and travel times by hand based on the closest feasible point (This is the case for some zipcodes in Valais and Grisons). Once the travel distances and -times are calculated, we aggregate them on a municipality level using the weights provided by the Swiss Federal Office of Statistics.

Appendix 3.G Heterogeneity with Respect to Sociodemographics

Table 3.G.1. Heterogeneity with Respect to Sociodemographics: UE Duration

	Pooled DiD	Female	Male	LPPS[a]	Blue Collar	White Collar	Marr.	Unmarr.	Low Edu.	Medium Edu.	High Edu.	Non-Swiss	Swiss
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)
DiD	14.472***	8.318*	17.069***	17.069***	24.620***	6.053	17.267***	11.197***	26.069***	12.606***	3.347	21.827***	8.282**
	(3.728)	(4.715)	(4.174)	(4.174)	(4.483)	(7.774)	(4.826)	(3.789)	(7.008)	(4.299)	(4.413)	(4.956)	(4.061)
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Outc. Mean	266.27	280.05	258.74	258.74	244.24	295.71	276.26	256.57	278.34	254.41	270.21	265.88	266.62
R^2	0.114	0.106	0.125	0.125	0.165	0.123	0.126	0.113	0.179	0.107	0.106	0.144	0.107
N	348’448	123’162	225’286	225’286	120’824	53’991	171’637	176’811	99’211	137’941	111’296	167’251	181’197

Notes: [a]: Low Paid Service Sector

This table shows the results from pooled Difference-in-Differences regressions, conditional on job seeker/unemployment spell characteristics. The outcome in all specifications is unemployment duration, top-coded at 520 days. In column (1), we display the results from our original Diff-in-Diff specification. Columns (2) and (3) contain the results for female and male job seekers, respectively. In columns (4) - (6), we display the results for White Collar-, Blue Collar-, and Low Paid Service Sector workers, respectively. Columns (7) and (8) present the results for married and unmarried workers, and columns (9) - (11) present the results for low education-, medium education- and high education workers, respectively. Note the definition of these categories below. Columns (12) and (13) present the results for non-Swiss citizens and Swiss Citizens, and columns (12) and (13) for German and French native speakers, respectively.

Low Education: Max. Compulsory High School, Medium Education: Max. Vocational Training, High Education: Max. College

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

Table 3.G.2. Heterogeneity with Respect to Sociodemographics: Post-UE Earnings

	Pooled DiD	Female	Male	LPPS[a]	Blue Collar	White Collar	Marr.	Unmarr.	Low Edu.	Medium Edu.	High Edu.	Non-Swiss	Swiss
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)
DiD	0.024***	0.044***	0.016	0.041***	0.005	0.025*	0.025**	0.027**	0.024	0.025**	0.022	0.002	0.035***
	(0.008)	(0.009)	(0.014)	(0.014)	(0.012)	(0.015)	(0.012)	(0.011)	(0.019)	(0.012)	(0.015)	(0.012)	(0.009)
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Outc. Mean	8.47	8.32	8.54	8.56	8.47	8.12	8.47	8.46	8.28	8.41	8.69	8.40	8.53
R^2	0.396	0.401	0.364	0.410	0.306	0.224	0.460	0.338	0.336	0.286	0.401	0.402	0.382
N	187’490	64’175	123’315	94’153	67’148	26’189	89’331	98’159	49’496	77’185	60’809	88’061	99’429

Notes: [a]: Low Paid Service Sector

This table shows the results from pooled Difference-in-Differences regressions, conditional on job seeker/unemployment spell characteristics. The outcome in all specifications is log earnings after leaving unemployment. In column (1), we display the results from our original Diff-in-Diff specification. Columns (2) and (3) contain the results for female and male job seekers, respectively. In columns (4) - (6), we display the results for White Collar-, Blue Collar-, and Low Paid Service Sector workers, respectively. Columns (7) and (8) present the results for married and unmarried workers, and columns (9) - (11) present the results for low education-, medium education- and high education workers, respectively. Note the definition of these categories below. Columns (12) and (13) present the results for non-Swiss citizens and Swiss Citizens, and columns (12) and (13) for German and French native speakers, respectively.

Low Education: Max. Compulsory High School, Medium Education: Max. Vocational Training, High Education: Max. College

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

References

Arni, Patrick, and Amelie Schiprowski. 2019. "Job Search Requirements, Effort Provision and Labor Market Outcomes." *Journal of Public Economics* 169: 65–88. DOI: https://doi.org/10.1016/j.jpubeco.2018.09.004. [145, 146]

Ashenfelter, Orley, David Ashmore, and Olivier Deschênes. 2005. "Do Unemployment Insurance Recipients Actively Seek Work? Evidence from Randomized Trials in Four U.S. States." *Journal of Econometrics* 125 (1): 53–75. DOI: https://doi.org/10.1016/j.jeconom.2004.04.003. [145]

Bandura, Albert. 1977. "Self-efficacy: Toward a unifying theory of behavioral change." *Psychological review* 84 (2): 191. DOI: https://doi.org/10.1037/0033-295X.84.2.191. [143]

Card, David, and Phillip B Levine. 2000. "Extended Benefits and the Duration of UI Spells: Evidence from the New Jersey Extended Benefit Program." *Journal of Public Economics* 78 (1): 107–38. DOI: https://doi.org/10.1016/S0047-2727(99)00113-9. [145]

Chetty, Raj. 2008. "Moral Hazard Versus Liquidity and Optimal Unemployment Insurance." *Journal of political Economy* 116 (2): 173–234. DOI: https://doi.org/10.1086/588585. [145]

Cheung, Maria, Johan Egebark, Anders Forslund, Lisa Laun, Magnus Rodin, and Johan Vikström. 2019. "Does Job Search Assistance Reduce Unemployment? Experimental Evidence on Displacement Effects and Mechanisms." Working paper. DOI: http://dx.doi.org/10.2139/ssrn.3515935. [146]

Crépon, Bruno, Esther Duflo, Marc Gurgand, Roland Rathelot, and Philippe Zamora. 2013. "Do Labor Market Policies have Displacement Effects? Evidence from a Clustered Randomized Experiment." *Quarterly Journal of Economics* 128 (2): 531–80. DOI: 10.1093/qje/qjt001. [146]

DellaVigna, Stefano, and M Daniele Paserman. 2005. "Job Search and Impatience." *Journal of Labor Economics* 23 (3): 527–88. DOI: https://doi.org/10.1086/430286. [143]

Gautier, Pieter, Paul Muller, Bas van der Klaauw, Michael Rosholm, and Michael Svarer. 2018. "Estimating Equilibrium Effects of Job Search Assistance." *Journal of Labor Economics* 36 (4): 1073–125. DOI: https://doi.org/10.1086/697513. [146]

Guan, Yanjun, Hong Deng, Jiaqing Sun, Yanan Wang, Zijun Cai, Lihui Ye, Ruchunyi Fu, Yang Wang, Shu Zhang, and Yuhui Li. 2013. "Career Adaptability, Job Search Self-Efficacy and Outcomes: A Three-Wave Investigation Among Chinese University Graduates." *Journal of Vocational Behavior* 83 (3): 561–70. DOI: https://doi.org/10.1016/j.jvb.2013.09.003. [143]

Johnson, Terry R., and Daniel H. Klepinger. 1994. "Experimental Evidence on Unemployment Insurance Work-Search Policies." *Journal of Human Resources* 29 (3): 695–717. URL: http://www.jstor.org/stable/146249. [145]

Johnston, Andrew, and Alexandre Mas. 2018. "Potential Unemployment Insurance Duration and Labor Supply: The Individual and Market-Level Response to a Benefit Cut." *Journal of Political Economy* 126(6): 2480–522. [146]

Klepinger, Daniel H., Terry R. Johnson, and Jutta M. Joesch. 2002. "Effects of Unemployment Insurance Work-Search Requirements: The Maryland Experiment." *ILR Review* 56 (1): 3–22. DOI: 10.1177/001979390205600101. [145]

Lalive, Rafael. 2008. "How Do Extended Benefits Affect Unemployment Duration? A Regression Discontinuity Approach." *Journal of Econometrics* 142 (2): 785–806. DOI: https://doi.org/10.1016/j.jeconom.2007.05.013. [145]

Lalive, Rafael, Camille Landais, and Josef Zweimüller. 2015. "Market Externalities of Large Unemployment Insurance Extension Programs." *American Economic Review* 105 (12): 3564–96. URL: http://www.jstor.org/stable/43821386. [146]

Landais, Camille, Pascal Michaillat, and Emmanuel Saez. 2010. “Optimal Unemployment Insurance Over The Business Cycle.” Working paper. Cambridge, USA: National Bureau of Economic Research. eprint: http://eprints.lse.ac.uk/35596/. [146]

Landais, Camille, Pascal Michaillat, and Emmanuel Saez. 2018. “A Macroeconomic Approach to Optimal Unemployment Insurance: Theory.” *American Economic Journal: Economic Policy* 10 (2): 152–81. DOI: 10.1257/pol.20150088. [146]

Manning, Alan. 2009. “You Can't Always Get What You Want: The Impact of the UK Jobseeker's Allowance.” *Labour Economics* 16 (3): 239–50. DOI: https://doi.org/10.1016/j.labeco.2008.09.005. [145]

McVicar, Duncan. 2008. “Job Search Monitoring Intensity, Unemployment Exit and Job Entry: Quasi-Experimental Evidence from the UK.” *Labour Economics* 15 (6): 1451–68. DOI: https://doi.org/10.1016/j.labeco.2008.02.002. [145]

McVicar, Duncan. 2010. “Does Job Search Monitoring Intensity Affect Unemployment? Evidence from Northern Ireland.” *Economica* 77 (306): 296–313. DOI: https://doi.org/10.1111/j.1468-0335.2008.00747.x. [145]

Meyer, Bruce D. 1995. “Lessons from the U.S. Unemployment Insurance Experiments.” *Journal of Economic Literature* 33 (1): 91–131. URL: http://www.jstor.org/stable/2728911. [145]

Mueller, Andreas I., Johannes Spinnewijn, and Giorgio Topa. 2021. “Job Seekers' Perceptions and Employment Prospects: Heterogeneity, Duration Dependence, and Bias.” *American Economic Review* 111 (1): 324–63. DOI: 10.1257/aer.20190808. [143]

Nekoei, Arash, and Andrea Weber. 2017. “Does Extending Unemployment Benefits Improve Job Quality?” *American Economic Review* 107 (2): 527–61. DOI: 10.1257/aer.20150528. [145, 158]

Schmieder, Johannes F., Till von Wachter, and Stefan Bender. 2012. “ The Effects of Extended Unemployment Insurance Over the Business Cycle: Evidence from Regression Discontinuity Estimates Over 20 Years.” *Quarterly Journal of Economics* 127 (2): 701–52. DOI: 10.1093/qje/qjs010. [145]

Schmieder, Johannes F., Till von Wachter, and Stefan Bender. 2016. “The Effect of Unemployment Benefits and Nonemployment Durations on Wages.” *American Economic Review* 106 (3): 739–77. DOI: 10.1257/aer.20141566. [143]

Spinnewijn, Johannes. 2015. “Unemployed but Optimistic: Optimal Insurance Design with Biased Beliefs.” *Journal of the European Economic Association* 13 (1): 130–67. DOI: 10.1111/jeea.12099. [143]

Van den Berg, Gerard J, and Bas Van der Klaauw. 2006. “Counseling and Monitoring of Unemployed Workers: Theory and Evidence From a Controlled Social Experiment.” *International Economic Review* 47 (3): 895–936. DOI: https://doi.org/10.1111/j.1468-2354.2006.00399.x. [145]

Van den Berg, Gerard J, and Johan Vikström. 2014. “Monitoring Job Offer Decisions, Punishments, Exit to Work, and Job Quality.” *Scandinavian Journal of Economics* 116 (2): 284–334. DOI: https://doi.org/10.1111/sjoe.12051. [145]

Venn, Danielle. 2012. “Eligibility Criteria for Unemployment Benefits.” (131): DOI: https://doi.org/10.1787/5k9h43kgkvr4-en. [145]

Zimmerman, Ryan D., Wendy R. Boswell, Abbie J. Shipp, Benjamin B. Dunford, and John W. Boudreau. 2012. “Explaining the Pathways Between Approach-Avoidance Personality Traits and Employees' Job Search Behavior.” *Journal of Management* 38 (5): 1450–75. DOI: 10.1177/0149206310396376. [143]

Zweimüller, Josef. 2018. “Unemployment Insurance and the Labor Market.” *Labour Economics* 53: 1–14. DOI: https://doi.org/10.1016/j.labeco.2018.06.003. [152]