Modelling the Non-Linear Relation between Natural Gas Consumption and Temperature: An Alternative Framework

Hooijmans, Mike

MSc Financial Economics, Erasmus School of Economics (ESE), Erasmus University Rotterdam

Abstract

In this paper, I present an alternative modelling strategy to examine the non-linear relation between temperature and natural gas consumption for 5 European countries. Previous research has typically aimed to capture this non-linearity through the use of heating degree days (HDDs) and cooling degree days (CDDs) variables. However, while this approach is widely disseminated in the literature, it has several apparent drawbacks, i.e. a priori identification of the threshold value, and the ad-hoc transition from warmer to cooler regimes. In this study I examine the potentiality of different non-linear models to both describe the behaviour of natural gas demand, and to identify and validate the value(s) of the temperature threshold(s). Among the models under consideration, the most deliberate specification is the logistic smooth transition regression (LSTR) model. In contrast with the HDDs and CDDs approach, this method allows for a posteriori determination of the threshold value(s), thereby providing a method to examine their validity. Moreover, the LSTR model is more able to describe the degree of smoothness and the qualitative behaviour of the demand response function for values close to the threshold value. Lastly, by using an analogue of the Akaike Information Criteria (AIC), I formally show that that the LSTR specification outperforms the conventional HDDs and CDDs approach for all countries.

Keywords: Natural Gas Consumption; Temperature; Heating Degree Days; Cooling Degree Days; Regime-Switching Model; Threshold Regression Model; Smooth Transition Regression Model

1 Introduction

Over the past two decades the markets for natural gas in Europe have undergone a profound restructuring process. Traditionally, natural gas markets in most developed countries were restrictively monopolistic and government-controlled. However, in the late 1990s the liberalisation and privatisation of natural gas markets emerged prominently on the European political agenda (Haase, 2008). This process has been pivotal in the convergence of regulatory regimes across Europe, but recently attention has shifted towards *other*, more pressing, energy challenges. In particular, growing worries with regard to global warming set climate change at the heart of the European energy debate[1].

At the Paris climate conference (COP21) in December 2015, 195 countries agreed to the "first-ever universally, legally binding global climate deal"[2]. In short, governments agreed on a long-term goal of keeping the increase in the global average temperature below 2 °C, as this level of global warming is generally judged to inflict dangerous climatological changes. However, the majority of climate change assessments have centred on the contributions of the energy sector to global warming (Amato et al., 2005). Conversely, relatively few studies explore the reverse implications of climatological factors on the energy sector. In line with these findings, this paper will focus on the relation between natural gas consumption and temperature[3].

The primary studies on the link between energy consumption and temperature consistently indicate that the relation is non-linear (Peirson and Henley, 1994, Li and Sailor, 1995, Al-Zayer and Al-Ibrahim, 1996, Sailor and Muñoz, 1997, Henley and Peirson, 1997, 1998). This non-linearity derives from the fact that both increases and decreases in temperature, relative to certain threshold values, can lead to an increase in energy demand. In particular, during the winter the expected link between energy consumption and temperature is expected to be negative, i.e. lower (colder) temperatures are expected to cause an increase in energy demand (heating effect). Vice versa, during the summer the expected link between energy consumption and temperature is expected to be positive, i.e. higher (warmer) temperatures are expected to cause an increased need for cooling leading to an increase in energy demand (cooling effect).

Prior research has typically aimed to capture this non-linearity through the use of heating degree days (HDDs) and cooling degree days (CDDs) variables (Al-Zayer and Al-Ibrahim, 1996, Sailor and Munõz,

[1] European Commission: A policy framework for climate and energy in the period from 2020 to 2030.

[2] European Commission: Climate Action, Paris Agreement 2015.

[3] Throughout this paper I use the terms consumption and demand interchangeably to describe the use of energy (and in particular, natural gas).

1997, Considine, 2000, Valor et al., 2001, Pardo et al., 2002, Amato et al., 2005, Eskeland and Mideksa, 2009, Petrick et al., 2010). Conventionally, these are defined as follows, $HDD = \max(T^* - T_t, 0)$, and $CDD = \max(T_t - T^*, 0)$, where T^* is an arbitrary threshold temperature. Each degree deviation from this predefined balance point is then counted as a degree-day. For instance, if the arbitrary threshold value is set to 18 °C, and the average temperature on day t is 23 °C, then this would result in 5 CDDs for that day. Vice versa, for the same threshold value, if the temperature on day t is 13 °C, this would result in 5 HDDs for that day[4].

Jointly, these two functions specify the number of days on which the temperature exceeds or falls short of the threshold value, and by how many degrees. This approach enables the construction of a parsimonious linear regression model (that contains HDDs and CDDs as independent variables), which maintains ease of interpretation while still capturing the non-linear behaviour of energy demand. Although this modelling strategy has performed reasonably well in the previous literature, it also has several apparent shortcomings. Firstly, the identification of heating and cooling degree days is based on an arbitrary threshold value, typically chosen to be approximately 18 °C (i.e. 65 °F), whose validity is not subjected to any type of check. Secondly, it is nontrivial whether a single threshold or a dual threshold should be considered. The former would suggest a sharp change in the behaviour of demand for temperatures close to the threshold value, i.e. as depicted in figure 1, whereas the latter assumes a smoother transition with an intermediate temperature range in which there is no appreciable change in consumption.

Moral-Carcedo and Vicéns-Otero (2005) acknowledge the gravity of these shortcomings, and propose a different methodological approach to study the relation between temperature and daily electricity demand in Spain over the period 1995-2003. They suggest a simple regime switching model to examine both the implied existence of different regimes (cold and warm states), as well as the transition from one to the other. The behaviour that they observe clearly delineates a *cold* and *hot* regime, as well as an intermediate range in which a gradual transition occurs from the former to the latter. They then employ a threshold regression model to further corroborate the preliminary evidence of a U-shaped demand response function. Building on these results, the authors propose a smooth transition regression model with a logistic transition function (LSTR), to adequately capture the U-shaped non-linear relation between temperature and Spanish electricity consumption.

[4] The threshold value T^* can be different for the HDDs and CDDs specification, thereby allowing for a temperature range in which energy demand is insensitive to temperature variations.

Figure 1: The theoretical relation between energy consumption and temperature.

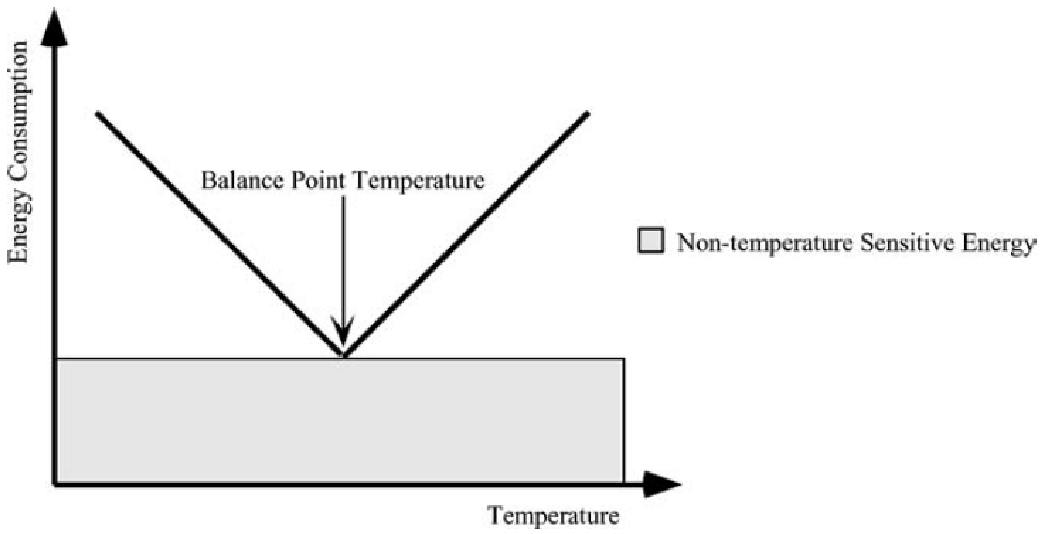

Notes: Portrays the V-shaped relation between energy consumption and temperature as theorized by the HDDs and CDDs approach.

Bessec and Fouquau (2008) extend the work of Moral-Carcedo and Vicéns-Otero (2005), as they examine the non-linear relation between electricity consumption and temperature for a panel of 15 European countries. They built on the theoretical work of González et al. (2005) and employ panel smooth transition regression models, implementing both logistic and exponential transition functions. Their main empirical findings confirm the non-linear U-shaped pattern of the relation, and they show that the non-linearity is more pronounced for warmer countries. Moreover, they endogenously estimate the threshold value to be 16.1 °C on average for the panel of European countries.

Lee and Chiu (2011) construct a similar research methodology, and develop a panel smooth transition regression model for 24 OECD countries. They further the analysis of Bessec and Fouquau (2008) by examining the sensitivity of electricity consumption to both electricity price and real income (, in addition to temperature). Similarly, they find evidence of a U-shaped relation between electricity consumption and temperature. Moreover, they endogenously estimate the threshold value of temperature to be approximately 12 °C (i.e. 53 °F), which is significantly lower than in Bessec and Fouquau (2008). For a more comprehensive review of the existing literature, I refer to Mideksa and Kallbekken (2010), Aufhammer and Mansur (2014), and Ranson, Morris and Kats-Rubin (2014).

1.1 The Relation between Natural Gas Consumption and Temperature

Auffhammer and Mansur (2014) observe that most of the prevailing research on energy demand sensitivity focuses on electricity consumption. Studies on other fuels are sparse, and in particular natural gas consumption is mostly neglected in the existing literature. De facto, to the best of my knowledge, there has not been any research which concentrates exclusively on the relation between natural gas consumption and temperature. However, similar to electricity demand, natural gas consumption is largely determined by the decisions of large groups of different economic agents (e.g. residential consumers, gas-fired power plants, commercial consumers etc.). Even assuming that these agents only make dichotomous decisions, and change their behaviour discretely, it seems unlikely that they all do this simultaneously. In other words, not every economic agent reacts similarly to temperature variations. This suggests that a U-shaped demand response function might likewise be more accurate in depicting the response of natural gas consumption to temperature variations.

Figure 2: Regime switching probability and temperature.

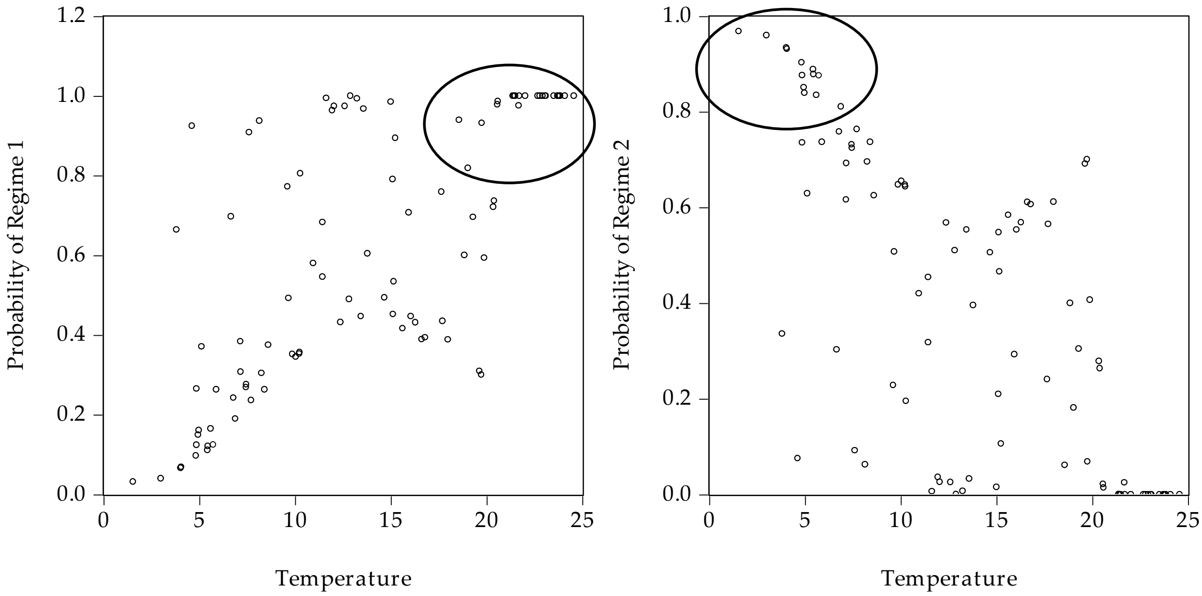

Notes: Depicts the results of a simple regime switching model with two states for Italy, without imposing any a priori relation between the switching probability and temperature. The left-hand figure portrays the probabilities of regime 1 plotted against temperature, while the right-hand figure shows the probabilities of regime 2. For a complete description of the underlying methodology I refer to section 2.

Figure 2 shows preliminary support for this line of reasoning by portraying the results of a simple regime switching model with two states for one of the sample countries, i.e. Italy. Without imposing any a priori relation between the switching probability and temperature, the link between regime and temperature still manifests itself clearly. In particular, regime one can be identified as a *warm* regime that exhibits a high probability of occurrence for temperatures above 20 °C. Whereas regime two corresponds with a *cold* regime, demonstrating a high probability of occurrence for temperatures below 10 °C.

Moreover, as temperature increases, the probability of regime one increases progressively without sudden jumps, i.e. it follows a gradual process with frequent *reverses* in the transition from one regime to the other. This behaviour corresponds with a smooth transition from winter to summer, as economic agents adjust their behaviour to prevailing temperature levels. These preliminary findings corroborate the intuition that similar to electricity demand, the response of natural gas consumption to temperature variations is not V-shaped as suggested by the HDDs and CDDs approach. Rather there is a smoother transition with an intermediate temperature range in which there is no appreciable change in consumption.

The main objective of this paper is to propose a modelling strategy that is able to adequately capture this non-linear effect of temperature on natural gas consumption. In the first step, I apply regime switching models and threshold regression models to demonstrate the shortcomings of the conventional HDDs and CDDs approach. Subsequently, I propose smooth transition regression (STR) models, an econometric method widely diffused in macro-economic research, to describe the response of natural gas consumption to temperature variations. In contrast with the HDDs and CDDs approach, this method allows for a posteriori determination of the threshold value, thereby providing a method to examine the validity of the threshold value. Furthermore, it captures more adequately the transition from warmer to cooler regimes, i.e. the response of natural gas demand to temperature changes for intermediate temperature ranges. To the best of my knowledge this is the first occasion where this type of model is used in the context of natural gas consumption.

This paper further adds to the existing literature by using data for five different European member states (France, Germany, Italy, the Netherlands, and the United Kingdom). This enables me to identify that both the shape and location parameter of the demand response curve differ meaningfully across countries. Therefore I argue that the assessment of demand sensitivity to temperature should be performed at the regional scale (for large countries such as the United States) or at country-level (as in the case of the European Union).

The remainder of this paper is organized in the following manner. In section 2 I provide an elaborate discussion of the methodology, which consists of a detailed portrayal of the different estimation techniques and model specifications. Section 3 serves as a description and analysis of the complete dataset. In section 4 I depict the empirical results and their implications. Finally, the main conclusions are presented and summarised in section 5.

2 Methodology

The preliminary results derived in the previous section suggest that the effect of temperature on natural gas consumption is non-linear. In particular, demand appears to increase both during warmer periods in the summer, as well as in colder periods during the winter. Conventionally, researchers have aimed to capture this non-linearity through the use of HDDs and CDDs variables. Even though this modelling strategy has performed reasonably well in the previous literature, it also has several apparent shortcomings as specified in the previous section.

In an attempt to overcome these difficulties I explore three alternative modelling strategies in the next sub-sections. These models aim to capture the non-linear relation between temperature and natural gas demand more adequately, focusing in particular on the switch in behaviour of natural gas demand for different temperature regimes. More specifically, in addition to the HDDs and CDDs variables approach, I employ a regime-switching regression model, a threshold regression model, and a smooth transition regression model. These approaches can be classified as alternating or switching models, in which different linear models are estimated for different states that occur in succession or alternate with each other. The characteristics of these models allow for a posteriori determination of the threshold value, and a smoother transition from colder to warmer temperature regimes.

2.1 Country-Level Energy Demand Sensitivities

The literature on energy demand sensitivities can broadly be divided in two competing strands which diverge in their methodological approach, i.e. individual time-series and panel studies. Similar to Amato et al. (2005), I argue that energy demand sensitivities with respect to climatic variables should be examined at the regional scale (for large countries such as the United States) or at country-level (as in the case of the European Union). In particular, as Boustead and Yaros (1994) show, energy infrastructures can differ significantly across countries and regions. For instance, energy systems might be different in terms of energy sources, efficiencies, age of transmission and distribution systems, and end-use technologies, thereby affecting the demand response curve. Secondly, as residential, commercial and industrial sectors exhibit different demand responses to temperature, the sectoral composition of a country's economy might significantly influence the sensitivity of natural gas demand (Amato et al., 2005). Therefore, I opt that it is sensible to estimate the forthcoming models individually for each country (i.e. contrasting with a panel approach).

2.2 HDDs and CDDs

In order to incorporate the non-linear dynamics of energy consumption in a linear model, the majority of the related literature has segmented the variation of temperature in two *new* variables (HDDs and CDDs). Thereby essentially segregating the behaviour of demand in colder and warmer regimes. These variables can more formally be defined as follows,

$$HDD = \max(T^* - T_t, 0) \tag{1}$$

$$CDD = \max(T_t - T^*, 0) \tag{2}$$

where T^* is an arbitrary threshold temperature. Each degree deviation from this predefined balance point is then counted as a degree-day. Although there is no consensus in the literature on the exact value of this threshold, it is typically chosen to be approximately 18 °C (i.e. 65 °F). However, most studies provide no considerate rationale for using this value. Therefore I opt to follow Bessec and Fouquau (2008), who in a panel study of 15 European countries endogenously estimate the threshold value to be 16.1 °C.

This approach enables me to construct a parsimonious linear model with HDDs and CDDs as explanatory variables,

$$FD_t = \alpha_0 + \beta_1 HDD_t + \beta_2 CDD_t + \varepsilon_t \tag{3}$$

where FD_t depicts filtered natural gas demand (see section 3 for the filtering methods), and ε_t is assumed to be an independently and identically normally distributed error term ($N(0, \sigma^2)$).

2.3 Regime-Switching Model

The link between natural gas consumption and temperature is expected to be different for colder temperatures, then for warmer. In particular, I will assume that this response can be represented by a straightforward linear model with two states,

$$FD_t = \alpha_j + \beta_j TMP_t + \varepsilon_t \tag{4}$$

where TMP_t portrays temperature at time t, $\varepsilon_t \sim N(0, \sigma_j^2)$, and J={1, 2} represents each of the two states (*cold* or *hot* temperatures). The parameter β_j can then be interpreted as the effect of temperature on natural gas demand when the respective state is either *cold* or *hot* (and α_j is an intercept that typically has no meaningful interpretation).

As a starting point of this regime-switching model I assume that natural gas demand alternates independently between either of the two aforementioned states. In other words, I do not assume any a priori relation between the switching probability and temperature. Now, presuming that the probability of being in a particular state is independent of past values, and denoting the information set at time t as ψ_{t-1}, I find that $P(j = 1|\, \psi_{t-1}) = P(j = 1) = p$, and therefore $P(j = 2) = (1 - p)$. Subsequently, assuming normality in the distribution term $\varepsilon_t \sim N(0, \sigma_j^2)$, the natural logarithm of the likelihood function can then be formally described as follows,

$$\ln L = \sum_{t=1}^{T} \ln \left[\sum_{j=1}^{2} f(FD_t|j,\, \psi_{t-1}) P(j|\, \psi_{t-1}) \right] \tag{5}$$

where,

$$f(FD_t|j,\, \psi_{t-1}) = \frac{1}{\sigma_j \sqrt{2\pi}} exp \left[-\frac{1}{2} \left(\frac{FD_t - \alpha_j - \beta_j TMP_t}{2\sigma_j} \right)^2 \right] \tag{6}$$

and $P(j|\, \psi_{t-1})$ depicts the probability that the model is in state j = {1, 2}, given the information set ψ_{t-1} at time t.

The resulting model can then be estimated by virtue of the expectation-maximization (EM) algorithm, as described in Quandt (1988). More specifically, assuming an initial estimation of the probability that observation t belongs to state j, one can maximize the natural logarithm of the likelihood function given this probability (the maximization-step). Subsequently, given the values of these parameters, the probability that observation t belongs to state j can be updated as follows,

$$P(J = j|\, \psi_{t-1}) = \frac{f(FD_t|J = j) P(J = j)}{\sum_{j=1}^{2} f(FD_t|J) P(J)} \tag{7}$$

yielding the Bayesian probability that the observed value of FD_t has been instigated by the model prevailing in state j (the expectation-step). This new probability estimate can then be used in the next recursive step, and the maximization and expectation steps can subsequently be repeated until sufficient convergence in the likelihood function is reached. For a more formal and complete portrayal of the EM algorithm I refer to the annex in the appendix of Moral-Carcedo and Vicéns-Otero (2005).

2.4 Two-Threshold Regression Model

A simple (one-) threshold regression model is conceptually similar to a linear regression with two regimes as depicted in the previous section. However, now the transition from the first regime to the second occurs ad-hoc as the threshold variable takes on different values relative to the threshold.

Similarly, a two-threshold regression model can be seen as a linear model with three regimes and two thresholds. Applied to the setting in this paper, this model can then be formally described as follows,

$$FD_t = \begin{cases} \alpha_1 + \beta_1 TMP_t + \varepsilon_t & TMP_t < \varphi_1 \\ \alpha_2 + \beta_2 TMP_t + \varepsilon_t & \varphi_1 \geq TMP_t < \varphi_2 \\ \alpha_3 + \beta_3 TMP_t + \varepsilon_t & TMP_t \geq \varphi_2 \end{cases} \tag{8}$$

where φ_1 and φ_2 are threshold values, such that $\varphi_2 > \varphi_1$, TMP_t is both an explanatory variable as well as the threshold variable, and ε_t is an independently and identically normally distributed error term ($N(0, \sigma^2)$).

The resulting model is estimated by a procedure of consecutive ordinary least squares (OLS) estimation in which the value of the thresholds is modified at each step (e.g. Tsay, 1989, 1998; Hansen, 2000). For each iteration the Akaike Information Criteria (AIC), or alternatively the Sum of Squared Residuals (SSR), is obtained, which can then be used to determine the optimal threshold values (i.e. those that provide the lowest AIC or SSR). However, one of the limitations of this approach is that the threshold values shouldn't be too close to the 0^{th} and 100^{th} percentile of the threshold variable, as there are too little observations in these extreme points to provide efficient estimates (Tsay, 1989). In order to circumvent this problem I restrict the threshold values to fluctuate between the 5^{th} and 95^{th} percentile.

2.5 Smooth Transition Regression Model

The preliminary findings in section 1 suggest that the transition from *cold* to *warm* regimes (and vice versa) occurs gradually, and not sudden. This behaviour can be described adequately by smooth transition regression (STR) models (Teräsvirta and Anderson, 1992, Granger and Teräsvirta, 1993, Teräsvirta, 1994, 1998), which capture the probability of the prevalence of one state through a continuous transition function. In particular, I propose the following application of this model to capture the response of natural gas consumption to temperature,

$$FD_t = \alpha_1 + \beta_1 TMP_t + (\alpha_2 + \beta_2 TMP_t) G(z_t; \gamma, c) + \varepsilon_t \tag{9}$$

where $G(z_t; \gamma, c)$ is a continuous transition function bounded by the interval $[0, 1]$, dependent on a transition variable z_t, a slope parameter γ, and a location parameter c. In the context of this study the transition variable z_t will be represented by the temperature level TMP_t. Furthermore, ε_t is assumed to be an independently and identically normally distributed error term ($N(0, \sigma^2)$).

2.5.1 Transition Function

The literature broadly distinguishes between two common types of transition functions, the logistic smooth transition (LSTR) and the exponential smooth transition (ESTR),

$$G(z_t; \gamma, c) = [1 + exp\{-\gamma(z_t - c)\}]^{-1} \quad (10)$$

$$G(z_t; \gamma, c) = 1 - exp\{-\gamma(z_t - c)^2\} \quad (11)$$

where (10) refers to the logistic specification, and (11) to the exponential. Although both functions are continuous, bounded by the interval [0, 1], and determined by the same parameters, there are some key differences. Figure 3 portrays these two transition functions for various values of the slope parameter γ (assuming c to be zero).

From figure 3 we observe that the logistic transition function has an S-shape. Hence, it could be used to describe the transition from a cooler regime to a warmer regime, with an intermediate range in which natural gas consumption is inelastic to temperature variations. Conversely, the exponential transition function is U-shaped, indicating that the two states correspond with a situation of extreme temperatures (both cold and warm), and one of intermediate temperature levels. Following Moral-Carcedo and Vicéns-Otero (2005) I opt for the logistic specification, as it seems counterintuitive to assume a priori that natural gas consumption responds similar to cold and hot temperature regimes.

Figure 3: Logistic and exponential smooth transition functions.

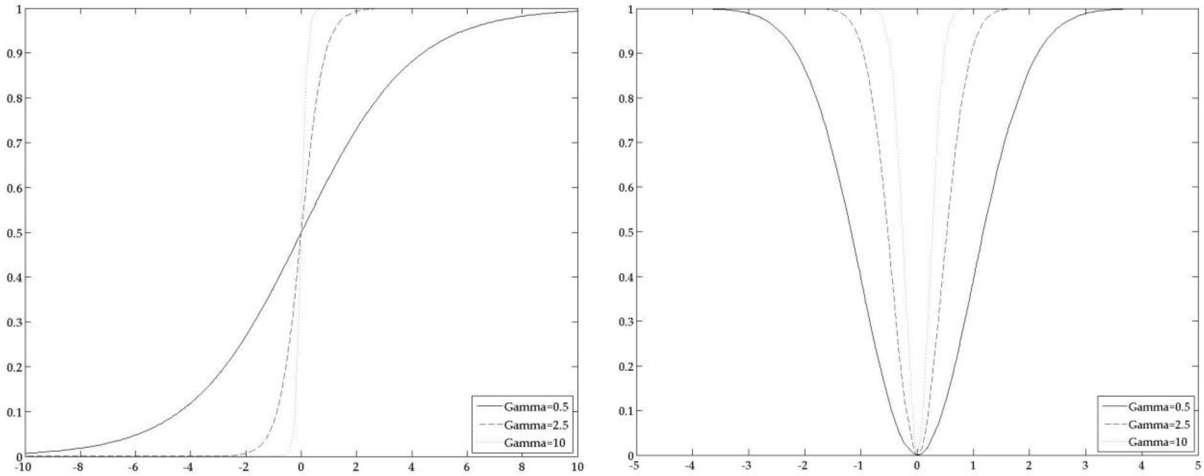

Notes: Figure 3 portrays examples of transition functions with on the left-hand side a logistic specification, and on the right-hand side an exponential specification. In both cases the location parameter c is assumed to be 0, and three different values for the shape parameter γ are depicted (i.e. 0.5, 2.5, and 10).

Moreover, figure 3 shows that the parameter γ, which determines the smoothness (i.e. speed of transition), can significantly affect the shape of the transition function. In particular, if the value of γ

tends to infinity the transition function becomes and indicator function, such that the model can be substituted with a threshold regression model with one threshold. Contrastingly, if γ tends to zero the transition function is constant, such that the model reduces to a standard linear regression model. In addition, note that the inclusion of the transition function allows for an infinite number of intermediate regimes. More specifically, the coefficients at each point in time can be seen as a weighted average of the values obtained in each of the two extreme regimes. Now, as the weights depend on the value of the transition function, this allows for a *continuum* of coefficient values, where each set is associated with a different value of the transition function.

2.5.2 Estimation and Linearity Test

The estimation of smooth transition regression models consists of several stages. Firstly, a linearity test should be conducted, which tests the LSTR model against a linear alternative. The null hypothesis for this linearity test can be expressed as $H_0 : \gamma = 0$, or alternatively $H_0 : \beta_1 = \beta_2$. However, in both circumstances there is an identification problem since the model is not identified under the null hypothesis. Namely, the model contains unidentified nuisance parameters (i.e. c, α_2, and β_2) under H_0, whose values do not affect the value of the log-likelihood (Davies, 1987). As a consequence, standard tests such as the Lagrange Multiplier, Likelihood Ratio, and Wald test, do not have their asymptotic distributions under the null hypothesis, rendering them invalid for consistent estimation of c, α_2, and β_2.

Luukkonen, Saikkonen, and Teräsvirta (1998) and Teräsvirta (1998) propose a solution to this problem by replacing the transition function with a first-order Taylor expansion around $\gamma = 0$. This yields the following equation for a logistic transition function,

$$FD_t = \theta_0 + \theta_1 TMP_t + \theta_2 TMP_t^2 + \varepsilon_t^* \tag{12}$$

such that testing linearity against the LSTR alternative reduces to $H_0 : \theta_2 = 0$. Now, if we denote the sum of squared residuals under the null hypothesis (i.e. the linear alternative) as SSR_0, and the sum of squared residuals under the alternative hypothesis (i.e. the LSTR model) as SSR_1, the LM-statistic is defined as follows,

$$LM_\chi = T\frac{(SSR_0 - SSR_1)}{SSR_0} \tag{13}$$

where T denotes the length of the times series, such that the LM-statistic has an approximate $\chi^2(1)$ distribution under the null hypothesis.

A similar testing method can be used to extend the LSTR model with one transition function, to a specification that allows for multiple transition functions. The procedure then consists of testing the null hypothesis of no remaining non-linearity. In particular, if the findings in (13) are such that linearity is rejected, one can test whether one transition function is appropriate ($H_0 : r = 1$), or whether there are at least two transition functions ($H_0 : r = 2$). Note that a STR model with two transition functions can be formally described as follows:

$$FD_t = \alpha_1 + \beta_1 TMP_t + (\alpha_2 + \beta_2 TMP_t) G_1(z_t; \gamma_1, c_1) + (\alpha_3 + \beta_3 TMP_t) G_2(z_t; \gamma_2, c_2) + \varepsilon_t \qquad (14)$$

Similarly, the procedure then entails replacing the second transition function with its first-order Taylor series expansion around $\gamma_2 = 0$, which yields the following model,

$$FD_t = \alpha_1 + \beta_1 TMP_t + (\alpha_2 + \beta_2 TMP_t) G_1(z_t; \gamma_1, c_1) + \theta_1 TMP_t^2 + \varepsilon_t^* \qquad (15)$$

The test of no remaining non-linearity is then simply defined as $H_0 : \theta_1 = 0$, and I can again construct the LM-statistic, where the sum of squared residuals under the null hypothesis (i.e. the specification with one transition function) is likewise denoted as SSR_0, and the sum of squared residuals of the alternative specification as SSR_1. A similar testing procedure can then be used to further extend the model to a specification that allows for more transition functions. If there is no remaining non-linearity the procedure ends, and the resulting model can then be estimated using non-linear least squares as proposed by Teräsvirta (1994, 1998).

3 Data

In this paper, I employ a dataset consisting of five European member states: France, Germany, Italy, the Netherlands, and the United Kingdom. These countries were primarily selected based on their volume of natural gas consumption, jointly accounting for approximately 70% of total natural gas consumption in the European Union[5]. The sample period extends from January 2008 to December 2015, and I gather monthly data for this period on natural gas consumption, population, production in total manufacturing, and temperature. Natural gas consumption is measured as gross inland consumption expressed in Terajoules, and data is obtained from Eurostat. Furthermore, production in total manufacturing is a seasonally unadjusted index obtained from the OECD database. The base year (originally 2010) is adjusted to 2008 to reflect the sample in this paper. Population data is similarly obtained from the OECD database, and following Bessec and Fouquau (2008) monthly data was attained by applying linear interpolation.

Moreover, temperature data is obtained from the high resolution gridded dataset constructed by Caesar et al. (2006)[6]. Daily observations on the minimum and maximum daily temperature are combined to form a daily average, which is then translated into HDDs and CDDs variables according to equation (1) and (2). In order to let the HDDs and CDDs variables coincide with the monthly frequency of the natural gas consumption data, the degree days are accumulated over time to provide monthly heating and cooling degree-day totals. The accuracy of the gridded dataset is crosschecked with the Cooperative Station Dataset published by the National Oceanic and Atmospheric Administration's National Climate Data Center. In particular, I gather monthly average temperature data expressed in degrees Celsius for 779 weather stations across the 15 member states. The individual time-series are then checked for completeness, spatial distribution and representativeness of population distribution, and then combined by taking the simple arithmetic mean to form 5 country-aggregates. In general, this method yields largely similar results to the high resolution gridded dataset of Caesar et al. (2006).

3.1 Descriptive Statistics

Table 1 portrays descriptive statistics on natural gas consumption and temperature for each of the five countries. It illustrates that the United Kingdom and Germany are the largest consumers of natural gas in absolute terms, whereas the Netherlands has (by far) the highest average consumption per capita. A possible explanation for this is a combination of the relatively small number of inhabitants in the

[5] Eurostat Energy Statistics – Supply, Transformation and Consumption (2014).
[6] MET Office Hadley Centre – Observation Datasets.

Netherlands, the excellent energy infrastructure, and the abundance of natural gas (i.e. there is a very large natural gas field located in the Northern part of the country). Moreover, table 1 depicts relatively large differences in minimum and maximum monthly natural gas consumption. This is mostly due to a slow-down of residential and commercial consumption in the summer as a result of reduced heating-requirements as well as through the reduction of industrial activity during the summer holiday, leading to lower consumption levels

Table 1: Descriptive statistics regarding natural gas consumption and temperature.

	Natural Gas Consumption				Temperature		
	Min.	Max.	Mean	Mean (Pc)	Min.	Max.	Mean
France	44,833	307,937	144,753	0.00222	2.83	22.36	12.87
Germany	126,201	490,862	277,864	0.00342	-3.34	19.81	9.06
Italy	115,557	411,023	237,380	0.00397	1.55	24.57	13.66
The Netherlands	67,073	244,402	130,064	0.00779	-2.02	20.34	10.22
United Kingdom	133,719	495,290	276,645	0.00436	0.64	18.51	10.50

Notes: Natural gas consumption refers to gross inland consumption expressed in Terajoules (TJ), and is obtained from Eurostat. The 4th column reflects natural gas consumption on a per capita basis, derived by dividing mean consumption with average population levels. Temperature is obtained from the high resolution gridded dataset constructed by Caesar et al. (2006), and reflects the simple arithmetic mean of the minimum and maximum temperature levels. Moreover, the table refers to the applicable sample period used in this paper, namely January 2008 to December 2015.

Furthermore, from Table 1 I observe that France and Italy are relatively *warm* countries, with an average monthly temperature close to 13 °C. For these countries I expect to observe a more pronounced cooling effect, and a more limited heating effect. The Netherlands and the United Kingdom on the other hand have a milder climate that is approximately three degrees colder on average, and Germany has the lowest average monthly temperature of the five sample countries. For these three *colder* countries I expect a more sizable heating effect, and a limited cooling effect as air-conditioning penetration is likely to be significantly lower.

3.2 Filter Approach

Since the aim of this paper is to examine the sensitivity of natural gas consumption to temperature, it is necessary to first eliminate the effect of unrelated trends and non-climatic influences. Following the previous literature I aim to filter out three components from natural gas demand. Firstly, the demographic trend, as the population over the whole area increased by more than 3% over the sample period this likely increased energy consumption. Secondly, the technological trend, as this leads to improving energy efficiency and different demand responses through for instance, increased air-conditioning, better home-insulation, and changes in appliance usage. Lastly, the monthly industrial seasonality, as a reduction in production during the summer period could for instance partially offset the theorized cooling effect.

In line with Moral-Carcedo and Vicéns-Otero (2005) and Bessec and Fouquau (2008) I firstly remove the demographic trend before applying subsequent alternative filters. This trend is filtered out simply by dividing natural gas consumption through the interpolated population levels. Then, the latter two components are eliminated by applying two alternate filters to the per capita natural gas consumption. These approaches are referred to as filter I, and II throughout the rest of this paper.

The first filter builds on the work of Moral-Carcedo and Vicéns-Otero (2005). They suggest that the filtered demand can be obtained as the residuals from an OLS regression of the unfiltered natural gas consumption per capita on a third degree time polynomial and a dummy for the month of August. As the aim of this dummy is to capture the decrease in production activity during the summer holiday, I extend it to include both July and August. This yields the following equation,

$$\mathrm{GC}_t = \beta_0 + \beta_1 t + \beta_2 t^2 + \beta_3 t^3 + \beta_4 D_{Summer,t} + FD_t^I \tag{16}$$

where $\mathrm{GC}_{i,t}$ represents the per capita natural gas consumption at time t, t is a simple time variable (0 for January 2008, 1 for February 2008, …), $D_{Summer,t}$ is a dummy variable which is 1 if observation t corresponds with July or August, and zero otherwise. Furthermore, non-significant terms are discarded from the model in a sequential manner, i.e. first the most insignificant term is discarded, and then the model is estimated again and re-evaluated for remaining insignificant terms.

The second filter is proposed by Bessec and Fouquau (2008) and follows a largely similar approach to Moral-Carcedo and Vicéns-Otero (2005). In particular, natural gas demand is regressed on a third degree time polynomial, and the summer dummy is replaced by a seasonally unadjusted production term. This approach might be more sensible in filtering out the seasonal effect of industrial activity as it might not be limited to the summer months. The model is then estimated as follows,

$$\mathrm{GC}_t = \beta_0 + \beta_1 t + \beta_2 t^2 + \beta_3 t^3 + \beta_4 Y_t + FD_t^{II} \tag{17}$$

where Y_t represents a seasonally unadjusted production term at time t. Again, following a similar approach as in (14), non-significant terms are sequentially discarded from the model.

In Appendix A1 and A2 I portray scatter plots of the filtered natural gas consumption plotted against temperature for each of the five countries under scrutiny. For illustration purposes a local polynomial kernel regression of order 2 is included with a Gaussian kernel between the two variables, i.e. the fitted line in the figures. As an illustrative example I consider the results obtained for Italy. From Appendix A1 I observe that the relation between natural gas consumption and temperature is clearly non-linear. For colder temperatures (i.e. below 15 degrees) demand decreases steadily as temperature increases, reflecting the decreased usage of heating appliances as temperature levels rise. Then between 15 and 20 degrees there is an intermediate range in which natural gas demand is inelastic with regard to temperature. Lastly, for temperature levels above this intermediate comfort zone, there is a positive relation between natural gas consumption and temperature. This positive link reflects the increased usage of cooling appliances, which indirectly, via electricity generation, affects natural gas demand. The results for France are largely similar to those of Italy, whereas in contrast, Germany, the Netherlands, and the United Kingdom all show a significantly less pronounced cooling effect. This likely reflects the colder climates in these countries, translating in a decreased need for air-conditioning and other cooling appliances as temperatures are seldom consistently at levels that necessitate cooling.

When comparing the results in appendix A1 to those in A2 I discern two important differences. Firstly, the dispersion in appendix A1 is larger than in A2, especially for higher temperatures. This could imply that the summer dummy employed in the first filtering approach might not fully capture the seasonal effect of production during the summer. Secondly, the cooling effect observed in appendix A1 disappears almost entirely for both France and Italy. Note however, that the relation between natural gas consumption and temperature remains non-linear, i.e. there is a gradual transition from the cooler regime to a temperature range in which demand is insensitive to temperature variations.

4 Empirical Results

This section presents the results of the previously described model specifications. First I discuss the findings of the conventional HDDs and CDDs modelling approach. Then I estimate the regime switching model and the two-threshold regression model, and draw preliminary inferences with regard to the shape of the demand response function. Lastly, I evaluate the findings of the LSTR specification, and formally compare them with the conventional HDDs and CDDs approach.

4.1 HDDs and CDDs Approach

Table 2 portrays the results of a simple regression of monthly filtered natural gas consumption on HDDs and CDDs. The analysis covers the sample period ranging from 2008 to 2015, where HDDs and CDDs are derived using a base temperature of 16.1°C. The standard errors are depicted in parentheses below the parameter estimates, and adjusted for heteroscedasticity and autocorrelation (using Newey-West standard errors). Furthermore, (I) denotes the model specification where filtered natural gas consumption is derived from equation (16) (i.e. Filter I), and vice versa (II) refers to filter II (as in (17)).

The constant term on the first row in table 2 captures the level of filtered monthly natural gas demand that is insensitive to temperature variations, i.e. the volume of filtered natural gas demanded at the balance point temperature.

Table 2: Regression results of HDDs and CDDs approach.

	France		Germany		Italy		The Netherlands		United Kingdom	
	(I)	(II)	(I)	(II)	(I)	(II)	(I)	(II)	(I)	(II)
Constant	-0.00259**	-0.00213**	-0.00167**	-0.00179**	-0.00121**	-0.00035**	0.00093**	0.0008**	-0.00077**	-0.00098**
	(0.00014)	(0.00011)	(0.00015)	(0.00015)	(0.00012)	(0.00014)	(0.00024)	(0.00021)	(0.00013)	(0.00011)
HDD	0.00027**	0.00025**	0.0002**	0.0002**	0.00033**	0.0003**	0.0005**	0.0005**	0.00028**	0.00031**
	(0.00002)	(0.00002)	(0.00001)	(0.00001)	(0.00001)	(0.00001)	(0.00003)	(0.00002)	(0.00002)	(0.00001)
CDD	0.00022**	-0.00005	0.00055**	0.00011	0.0002**	0.00002	0.00058**	0.00015*	0.00079**	0.00011
	(0.00004)	(0.00004)	(0.00008)	(0.00007)	(0.00002)	(0.00003)	(0.0001)	(0.00008)	(0.00018)	(0.00009)
R^2	0.813	0.8639	0.7719	0.8452	0.885	0.8979	0.8698	0.9011	0.7752	0.9135
SSR	1.79E-05	1.83E-05	2.38E-05	2.23E-05	1.69E-05	1.86E-05	7.46E-05	6.58E-05	2.94E-05	1.47E-05

Notes: Table 2 depicts regression results of equation (3). Column 1 for each country refers to the first filtering approach, and column 2 to the second filtering approach. Moreover, the analysis covers the sample period 2008-2015, and standard errors are denoted in parentheses and adjusted for heteroscedasticity and autocorrelation using Newey-West standard errors. In the last two rows I portray two commonly used evaluation techniques, namely the coefficient of determination as well as the Sum of Squared Residuals. Note that * = significant at 10 percent level, ** = significant at 5 percent level.

For France, Germany, Italy, and the United Kingdom the constant term is negative. This corresponds with the findings in Appendix A1 and A2, which illustrate that the balance point temperature coincides with a negative rate of filtered natural gas consumption. Conversely, the Netherlands has a positive constant term. This is likely due to the anomalously high level of per capita natural gas consumption in the Netherlands.

Table 2 further illustrates that monthly HDDs have a positive and significant effect on filtered natural gas consumption for all countries and filtering approaches (i.e. at conventional confidence levels). The positive sign for all parameters confirms the intuition that an increase in HDDs will lead to an increase in filtered natural gas demand. Conversely, the results for CDDs are less unanimous, differing significantly across countries and filtering approaches. In particular, the results for CDDs appear to be significant and positive for the first filtering approach, and insignificant for the second filtering approach. This confirms the findings in section 3.2 and appendix A1 and A2, which suggest that there is no distinctive cooling effect after applying the second filter, but rather a temperature range in which demand is largely insensitive to temperature variations.

Moreover, the evaluation criteria at the bottom of table 2 indicate that the HDDs and CDDs modelling approach explains a sizable proportion of the variance in the dependent variable. In particular, the coefficient of determination ranges between 0.77 and 0.92. These relatively high values are consistent with previous research on the relation between temperature and electricity/natural gas consumption (e.g. Pardo et al., 2002, Amato et al., 2005, Eskeland and Mideksa, 2009, etc.). In addition, the sum of squared residuals is relatively small, which corresponds with the small absolute values of the dependent variable, rather than implying an anomalously good fit. In section 4.4.3 I will use the sum of squared residuals in order to formally compare this approach with the LSTR specification.

4.2 Regime-Switching Model

The results in the previous section are largely consistent with prior research and confirm the cogency of HDDs and CDDs as explanatory variables. However, as described earlier this approach assumes a sharp change in the behaviour of demand for temperatures close to the threshold value. Even assuming that economic agents only make dichotomous decisions, and change their behaviour discretely, it seems unlikely that they all do this simultaneously. In other words, not every economic agent reacts similarly to temperature variations. This suggests the potential existence of a smoother transition with an intermediate temperature range in which there is no appreciable change in consumption, i.e. a U-shaped demand response function. In this section I aim to provide support for this theorization by further expanding the preliminary analysis of the regime-switching model described in section 1.

Appendix A3 portrays the results of the simple regime switching model with two states for filter I. Note that I assume that natural gas demand alternates independently between either of the two aforementioned states. In other words, I do not assume any a priori relation between the switching probability and temperature. In general, I find that regime one can be identified as a warm regime that exhibits a high probability of occurrence for warmer temperatures, i.e. the regime is prevalent across all countries for temperature levels close to their respective upper percentiles. Subsequently, for lower temperature levels, the dynamics of the switching probability are less unanimous across countries. For Italy I observe a gradual decrease in the probability of regime one as temperature decreases. This behaviour corresponds with a smooth transition from winter to summer, as economic agents adjust their behaviour to prevailing temperature levels.

On the other hand, for France, Germany, The Netherlands, and the United Kingdom there appears to be a more immediate decrease in the transition probability. In particular, these countries show a concentration of low switching probabilities (between 0 and 0.2) in the range from 12 °C to 16 °C, which then gradually increases as temperature levels further decrease. This suggests the potential existence of three regimes, i.e. one regime for temperature levels in the upper percentiles, a second for temperatures in the comfort zone (between 12 °C and 16 °C) where demand is likely to be inelastic to temperature variations, and a third for lower temperatures. These preliminary findings corroborate the intuition that the response of natural gas consumption to temperature variations is not V-shaped as suggested by the HDDs and CDDs approach. Rather there is a smoother transition with an intermediate temperature range in which there is no appreciable change in consumption.

Appendix A4 depicts the results of a simple regime switching model with two states for filter II. When compared to appendix A3, I find that the findings are largely dissimilar for Italy, the Netherlands, and the United Kingdom. In particular, for these three countries there appears to be a clear heating effect for temperature levels in the lower percentiles, whereas regime 1 prevails for the rest of the temperature range. For both Germany and France the findings are more comparable to appendix A3, i.e. they show a prevalent warm regime, and a more gradual transition in which probabilities decrease steadily as temperature decreases.

4.3 Two-Threshold Regression Model

The results of the analysis in the previous section suggest that the transition between the colder and warmer regime occurs gradually (i.e. for most countries). In particular, appendix A3 illustrates that there are potentially even three regimes for filter I, i.e. cold, intermediate, and warm. In this section I further corroborate these findings by employing a two-threshold regression model of filtered per capita natural gas consumption on temperature (assuming temperature itself to be the threshold variable). Note, that a single threshold type model is conceptually similar to a linear regression with two regimes as depicted in the previous section, except now the transition occurs ad-hoc and is based on an observable threshold variable. Similarly, a two-threshold regression model can be seen as a linear model with three regimes and two thresholds.

Appendix A5 to A9 depict the results of a two-threshold regression model for filter I. In particular, I portray the Akaike Information Criteria for different values of the two temperature thresholds. Since the results appear to be visually similar in the 3D plot (upper-left figure), I further delineate the plot in three 2D figures (respectively the upper-right, and lower left and right figures). As a starting point I consider the upper-right figure, from which I observe that across all countries the AIC spans a broad range for lower temperature levels. The range then gradually grows denser as temperature levels increase towards the upper percentiles. Intuitively, it might seem that as temperature increases the AIC decreases, i.e. the model appears to perform better. However, this increase in density mostly represents the fact that there are less combinations of the two threshold values as the first threshold increases (since necessarily φ_2 has to be larger than φ_1).

Therefore, the upper-right figure should be analysed conjointly with the lower-left and right figure. From this we observe that the lowest AIC values are achieved for relatively high values of the second threshold. In particular, for Germany, the Netherlands, and the United Kingdom for temperatures between 14 °C and 18 °C, and for France and Italy for temperatures between 18 °C and 22 °C. Now, if I consider these temperature ranges for the second threshold when examining the lower-right figure, I find that very similar AIC values are achieved for different values of the first threshold. This is indicative of a smooth transition, since if there was a more sudden transition the values of the thresholds that minimize the AIC would be more defined (i.e. the dark blue regions in the lower-right figures would stand out more).

In addition, table 3 illustrates the results of the two-threshold regression model for the threshold values that minimize the AIC. The findings for filter I confirm the previously observed range for the second threshold, and also support the broader dispersion found for the first threshold. When comparing these results with those of filter II I observe two important differences. Firstly, there appears to be less dispersion in the first threshold for the second filtering approach, as it centres around 10-11 °C. Secondly, the values for the second threshold are significantly lower than those found for the first filtering approach. This likely reflects the absence of a cooling effect in the second filtering approach (i.e. appendix A2 depicts a flattening of the curve for higher temperature levels), suggesting there is a cold regime for temperatures below the 10-11 °C range. Followed by a smooth transition for temperatures between the first and second threshold, to a temperature range in which demand is largely insensitive to temperature variations (i.e. for values above the second threshold).

Table 3: Estimation results of the two-threshold regression model.

	Filter I			Filter II		
	Threshold 1	Threshold 2	AIC	Threshold 1	Threshold 2	AIC
France	10 °C	19 °C	-12.78	14.25 °C	16.25 °C	-12.65
Germany	10.25 °C	17 °C	-12.52	10.25 °C	15 °C	-12.42
Italy	11.5 °C	21 °C	-12.84	11 °C	16.5 °C	-12.60
The Netherlands	5.75 °C	15 °C	-11.23	10 °C	11 °C	-11.33
United Kingdom	15.25 °C	16.5 °C	-12.47	11.5 °C	14 °C	-12.88

Notes: Depicts the results of the two-threshold regression model for the threshold values that minimize the AIC. The model is estimated by sequentially performing Ordinary Least Squares for different values of the thresholds. For a more complete description of the methodology I refer to section 2.4.

4.4 Logistic Smooth Transition Regression

The analysis until now suggests that the transition between warmer and cooler regimes occurs gradually rather than sudden as implied by the HDDs and CDDs approach. In particular, for the first filtering approach the response of per capita natural gas consumption to temperature appears to follow a U-shape (with a more pronounced cooling effect for Italy and France). Conversely, for the second filtering approach there appears to be no observable cooling effect, i.e. rather there is an extended temperature range in which demand is largely inelastic to temperature variations (smooth L-curve). Both types of behaviour can be adequately captured by smooth transition regression models as proposed by Teräsvirta and Anderson (1992).

4.4.1 Linearity tests

The estimation of smooth transition regression models consists of several stages. Firstly, a linearity test should be conducted, which tests the LSTR model against a linear alternative. Following, Luukkonen, Saikkonen, and Teräsvirta (1998) and Teräsvirta (1998) I estimate equation (12) and test $H_0 : \theta_2 = 0$, through an LM test. Subsequently, I follow a sequential testing procedure to test for remaining non-linearity and the existence of additional transition functions. In table 4 I depict the results of this test for the different countries and filtering approaches. The LM-statistic is $\chi^2(1)$ distributed under the null hypothesis, and the corresponding p-values are denoted in parentheses.

Table 4: Test for remaining non-linearity in the LSTR model specification.

	Filter I				
Transition Function	France	Germany	Italy	The Netherlands	United Kingdom
$H_0: r = 0$ vs $H_1: r = 1$	35.24	25.29	49.21	26.38	31.45
	(0.0000)	(0.0000)	(0.0000)	(0.0000)	(0.0000)
$H_0: r = 1$ vs $H_1: r = 2$	3.92	1.63	4.72	2.91	4.48
	(0.2701)	(0.6527)	(0.1826)	(0.4057)	(0.2139)
	Filter II				
Transition Function	France	Germany	Italy	The Netherlands	United Kingdom
$H_0: r = 0$ vs $H_1: r = 1$	13.98	9.87	36.47	16.30	15.26
	(0.0029)	(0.0197)	(0.0000)	(0.0001)	(0.0016)
$H_0: r = 1$ vs $H_1: r = 2$	2.91	2.78	0.47	2.41	1.55
	(0.4051)	(0.4262)	(0.9245)	(0.4913)	(0.6704)

Notes: Depicts the results of a LM-type test for no remaining non-linearity. The test statistic is approximately $\chi^2(1)$ distributed, and the p-values are denoted in parentheses below. For a more formal description of the test statistic I refer to sub-section 2.5.2.

Table 4 illustrates that the null hypothesis of linearity is convincingly rejected for all filtering approaches and across all countries. In general, the rejection appears to be stronger for the first filtering approach than for the second. Similarly, the rejection is also more robust for warmer countries than for colder countries. This likely results from the more pronounced cooling effect observed for the first filter, as well as for the subset of warmer countries (i.e. France and Italy), as illustrated in appendix A1 and A2. Moreover, the specification tests of no remaining non-linearity shows that for all countries and filtering approaches, one transition function is optimal. These results confirm the findings of Bessec and Fouquau (2008), who show that a small number of regimes is sufficient to capture the non-linearity of energy demand.

4.4.2 Estimation Results

Table 5 depicts the parameter estimates of the different LSTR model specifications. Before interpreting the results, recall that the LSTR specification allows for an evaluation of the effect of temperature on natural gas consumption given a certain temperature level. Therefore the coefficients in table 5 (i.e. α_1 and β_1) can be different from the parameter estimates in the extreme regimes (Bessec and Fouquau, 2008). As a results, it is generally preferred to only interpret the sign of these parameters, which indicate an increase or decrease in the coefficients depending on the temperature level.

Firstly, examining the values of α_1 and β_1 I observe that they are all significantly different from zero at conventional confidence levels. Moreover, corresponding with appendix A1 and A2 I find that α_1 is positive across all countries and filtering approaches, and β_1 is generally negative. The negative coefficients for β_1 are indicative of a heating effect in the winter, i.e. a decrease in temperature results in an increase in natural gas consumption for heating purposes. The only exception for which β_1 is not negative, is for the Netherlands (in particular, for filter II). However, this likely results from a combination of the rather large negative value for α_1, as well as the very small value for the slope parameter γ (recall that as γ tends to zero the LSTR specification reduces to a linear regression model).

The coefficient estimates for α_2 and β_2 are less unanimous across countries and filtering approaches. More specifically, α_2 and β_2 are generally not significant for the first filtering approach. A possible explanation for this is that the cooling effect observed in appendix A1 is not significant, i.e. rather there is a range of warmer temperatures in which natural gas demand is largely inelastic to temperature variations (corresponding with a smooth L-curve). Alternatively, I could argue from a different perspective that there is in fact a cooling effect, but that it is largely captured by the positive coefficient for α_2.

Table 5: Estimation results for the LSTR model specification.

	France		Germany		Italy		The Netherlands		United Kingdom	
	(I)	(II)	(I)	(II)	(I)	(II)	(I)	(II)	(I)	(II)
α_1	0.00148**	0.00174**	0.00146**	0.00142**	0.00391**	0.00429**	0.00896**	4.74031**	0.00393**	0.00411**
	(0.00013)	(0.00016)	(0.00008)	(0.00009)	(0.00015)	(0.00044)	(0.00021)	(0.15218)	(0.00013)	(0.00012)
β_1	-0.00023**	-0.00023**	-0.00019**	-0.00019**	-0.0003**	-0.00023*	-0.0005**	0.02903**	-0.00031**	-0.00033**
	(0.00001)	(0.00002)	(0.00001)	(0.00001)	(0.00002)	(0.00014)	(0.00003)	(0.01073)	(0.00001)	(0.00002)
α_2	0.00384	-0.00334**	0.00467	-0.00662**	0.01326	-0.00381**	-0.00339	-10.58935**	0.00194	-0.0023**
	(0.00781)	(0.00088)	(0.01115)	(0.00298)	(0.03481)	(0.00077)	(0.01415)	(0.07065)	(0.00411)	(0.00103)
β_2	-0.00008	0.0002**	-0.00013	0.0004**	-0.00035	0.0002*	0.00036	0.00748**	0.00000	0.00018**
	(0.00035)	(0.00005)	(0.00057)	(0.00016)	(0.00135)	(0.00011)	(0.00070)	(0.00108)	(0.00023)	(6.58659)
γ	1.74**	10	1.64**	10.04**	0.62**	0.51	1.15	0.01**	4.67**	10
	(0.69010)	(56.2918)	(0.71042)	(0.00967)	(0.16352)	(0.33111)	(0.91760)	(0.00414)	(2.26989)	(44.4819)
c	19.29**	15.73**	17.04**	16.42**	21.91**	11.76**	16.25**	16.96**	15.56**	13.88**
	(0.51193)	(0.98949)	(0.82478)	(2.08277)	(2.81483)	(2.88837)	(1.98400)	(0.00500)	(0.13175)	(0.58931)
SSR	1.39E-05	1.77E-05	1.85E-05	2.16E-05	1.52E-05	1.76E-05	6.97E-05	6.38E-05	1.93E-05	1.34E-05

Notes: Depicts the parameter estimates for the LSTR specification across all countries and filtering approaches. The standard errors are in parentheses below the coefficients, and corrected for heteroscedasticity. For each model the optimum number of transition functions was determined using a sequential testing procedure (see table 4). The parameter notation refers to equation (9), i.e. $FD_t = \alpha_1 + \beta_1 TMP_t + (\alpha_2 + \beta_2 TMP_t)G(z_t; \gamma, c) + \varepsilon_t$, where γ is a slope parameter, and c the location parameter. Note that the parameter estimates cannot be directly interpreted as elasticities.

For the second filtering approach I generally observe negative and significant values for α_2, and positive and significant values for β_2. The former likely reflects the lower level of per capita natural gas consumption for higher temperatures (as observed in appendix A2), whereas the latter term can be seen as a counter-effect that reduces the decrease resulting from the negative value of α_2.

Continuing with the slope parameter γ I find that the estimated value is rather small in general (i.e. between 0.5 and 2). Recall that as the slope parameter becomes smaller that the transition becomes smoother, whereas if it tends to infinity the transition becomes more abrupt as the transition function approximates an indicator function. For the first filtering approach I find that the slope parameter consistently fluctuates around one. Contrastingly, for the second filtering approach the results are more disparate. In particular, for France, Germany, and the United Kingdom the slope parameter approximates 10, indicating an abrupt transition. Whereas, for Italy and the Netherlands, the slope parameter is considerably smaller, implying a more gradual transition.

The location parameter c indicates the temperature level at which the transition function reaches an inflection point. For the first filtering approach I observe that the location parameter is congruent with the climatic condition of a country, i.e. c is larger for Italy and France, then for the other countries. Moreover, I observe that the location parameter is significantly different from the threshold value of 18.3 °C that is typically used in the literature. In particular, for the warmer countries the location parameter tends to be above this threshold value, whereas for Germany, the Netherlands, and the United Kingdom it is significantly lower. Observe that a threshold value of 16.1 °C for European countries, as suggested by Bessec and Fouquau (2008), would be reasonable for Germany, the Netherlands, and the United Kingdom. These findings illustrate the advantages of the methodological approach presented in this paper. Namely, not only is the location parameter estimated rather than imposed a priori, but table 5 also illustrates that there are intricate differences between countries.

4.4.3 Model Comparison

In order to formally compare the difference between the conventional HDDs and CDDs approach, and the LSTR model, I use an analogue of the Aikaike Information Criterion. In particular, following Burnham and Anderson (1998) I use the sum of squared residuals to compute the following,

$$AIC = t \ln\left({}^{SSR}/_{t}\right) + 2(p+1) \quad (18)$$

where t is the number of observations (i.e. 96 in this study), and p the number of parameters estimated. Then in order to examine whether two models (non-nested, but with the same dependent variable) are significantly different one can look at the difference in the AIC value. Burnham and Anderson

(1998, p. 123) suggest that a difference of around 4 to 7 roughly corresponds with a "95% confidence level".

In table 6 I portray the AIC for the HDDs and CDDs approach and the LSTR specification, as well as the difference between these models. For the first filtering approach the differences are consistently larger than 4, implying that the LSTR specification is superior to the HDDs and CDDs approach. Especially, for France, Germany and the United Kingdom the LSTR model strongly outperforms the conventional approach. For the second filtering approach the results are less definite. However, observe that the LSTR specification is always better than the conventional approach, albeit not significantly (i.e. with the exception of the United Kingdom).

Table 6: Testing the performance of the LSTR specification.

	France		Germany		Italy		The Netherlands		United Kingdom	
	(I)	(II)	(I)	(II)	(I)	(II)	(I)	(II)	(I)	(II)
AIC_{HDD}	-1479.5	-1477.4	-1452.2	-1458.4	-1485	-1475.8	-1342.5	-1354.6	-1431.9	-1498.4
AIC_{LSTR}	-1502.1	-1478.8	-1474.2	-1459.5	-1493.5	-1479	-1347	-1355.4	-1470.4	-1505.1
$\Delta\, AIC$	-22.61	-1.41	-22.07	-1.07	-8.47	-3.12	-4.55	-0.89	-38.47	-6.65

Notes: Describes the results of an analogue of the AIC as portrayed in equation (18). The number of parameters for the HDD and CDD approach is 3, and for the LSTR specification 4. The last column denotes the difference between the HDDs and CDDs approach with the LSTR specification, where a difference of between 4 and 7 roughly corresponds with a 95% confidence interval (Burnham and Anderson, 1998).

5 Conclusion

In this paper, I present an alternative modelling strategy to explore the non-linear relation between temperature and natural gas consumption. Previous research has typically aimed to capture this non-linearity through the use of heating degree days (HDDs) and cooling degree days (CDDs). Jointly, these two variables specify the number of days on which the temperature exceeds or falls short of the threshold value, and by how many degrees. However, while this approach is widely disseminated in the literature, it has several apparent drawbacks, i.e. a priori identification of the threshold value, and the ad-hoc transition from warmer to cooler regimes. This study examines the potentiality of different non-linear models to both describe the behaviour of natural gas demand, and to identify and validate the values of the temperature thresholds.

Among the models under consideration, the preferred specification is the logistic smooth transition regression (LSTR) model. In contrast with the HDDs and CDDs approach, this method allows for a posteriori determination of the threshold value, thereby providing a method to examine the validity of the threshold value(s). In particular, I show that the location parameter is generally significantly different from the threshold value of 18.3 °C that is typically used in the literature. In addition, I illustrate that both the shape and location parameter of the demand response curve differ meaningfully across countries. Therefore I argue that the assessment of demand sensitivity to temperature should be performed at the regional scale (for large countries such as the United States) or at country-level (as in the case of the European Union).

Furthermore, I show that the LSTR specification captures more adequately the transition from warmer to cooler regimes, i.e. the response of natural gas demand to temperature changes for intermediate temperature ranges. In particular, the LSTR model is more able to describe the degree of smoothness and the qualitative behaviour of the demand response function for values close to the threshold value. Ultimately, by using an analogue of the Akaike Information Criteria (AIC), I formally show that that the LSTR specification outperforms the conventional HDDs and CDDs approach for all countries.

5.1 Discussion

A limitation of the analysis I propose in this paper is that I do not distinguish between residential, commercial and industrial natural gas consumption. Dissecting aggregate natural gas consumption in these different categories is informative as they likely exhibit different demand responses with regard to temperature. Similarly, this study examines the temperature sensitivity of natural gas consumption at the national level. By focusing on such *large* geographical areas, I might forego to explicitly account for within-country differences in temperature, energy infrastructure, and sectoral composition. It is evident that a further decomposition in sectors and regions would be more relevant, however obtaining disaggregated data with a monthly frequency for this set of countries was simply unfeasible.

Moreover, in this paper I forego to explicitly test for the potentiality of different types of transition functions, i.e. I reject the notion of an exponential smooth transition function beforehand based on its shape, without applying any formal test. To examine the robustness of my choice, I can follow Escribano and Jorda (1999) who develop several Lagrange Multiplier (LM) type tests to choose between logistic and exponential smooth transition functions. Furthermore, in this study I propose the use of an analogue of the AIC in order to formally compare the difference between the conventional HDDs and CDDs approach, and the LSTR model. While Burnham and Anderson (1998) suggest that a difference of around 4 to 7 roughly corresponds with a "95% confidence level", they do not provide any rigorous statistical evidence to support this claim. Hence, it might be sensible to perform alternative tests for non-nested models to examine the robustness of the AIC method, and the findings in this paper.

6 References

Al-Zayer, J., & Al-Ibrahim, A. A. (1996). Modelling the impact of temperature on electricity consumption in the eastern province of Saudi Arabia. Journal of Forecasting, 15(2), 97-106.

Amato, A. D., Ruth, M., Kirshen, P., & Horwitz, J. (2005). Regional energy demand responses to climate change: Methodology and application to the Commonwealth of Massachusetts. Climatic Change, 71(1-2), 175-201.

Auffhammer, M., & Mansur, E. T. (2014). Measuring climatic impacts on energy consumption: A review of the empirical literature. Energy Economics, 46, 522-530.

Bessec, M., & Fouquau, J. (2008). The non-linear link between electricity consumption and temperature in Europe: a threshold panel approach. Energy Economics, 30(5), 2705-2721.

Boustead, I., & Yaros, B. R. (1994). Electricity supply industry in North America. Resources, conservation and recycling, 12(3-4), 121-134.

Burnham, K. P., & Anderson, D. R. (1998). Model selection and inference: a practical information-theoretic approach Springer-Verlag. *New York*.

Caesar, J., Alexander, L., & Vose, R. (2006). Large-scale changes in observed daily maximum and minimum temperatures: Creation and analysis of a new gridded data set. Journal of Geophysical Research: Atmospheres, 111(D5).

Considine, T. J. (2000). The impacts of weather variations on energy demand and carbon emissions. Resource and Energy Economics, 22(4), 295-314.

Davies, R. B. (1987). Hypothesis testing when a nuisance parameter is present only under the alternative. Biometrika, 74(1), 33-43.

Escribano, A., & Jorda, O. (1999). Improved testing and specification of smooth transition regression models. In *Nonlinear time series analysis of economic and financial data* (pp. 289-319). Springer US.

Eskeland, G. S., & Mideksa, T. K. (2009). Climate Change Adaptation and Residential Electricity Demand. Working Paper

European Commission: Climate Action, Paris Agreement 2015 – Accessed on the 12th of May 2016 at: http://ec.europa.eu/clima/policies/international/negotiations/paris/index_en.htm.

González, A., Teräsvirta, T., & Dijk, D. V. (2005). Panel smooth transition regression models (No. 604). SSE/EFI Working Paper Series in Economics and Finance.

Granger, C. W., & Terasvirta, T. (1993). Modelling non-linear economic relationships. OUP Catalogue.

Haase, N. (2008). European gas market liberalization: Are regulatory regimes moving towards convergence? Oxford Institute for Energy Studies. NG 24.

Hansen, B. E. (2000). Sample splitting and threshold estimation. Econometrica,68(3), 575-603.

Henley, A., & Peirson, J. (1997). Non-Linearities in Electricity Demand and Temperature: Parametric Versus Non-Parametric Methods. Oxford Bulletin of Economics and Statistics, 59(1), 149-162.

Henley, A., & Peirson, J. (1998). Residential energy demand and the interaction of price and temperature: British experimental evidence. Energy Economics, 20(2), 157-171.

Lee, C. C., & Chiu, Y. B. (2011). Electricity demand elasticity's and temperature: Evidence from panel smooth transition regression with instrumental variable approach. Energy Economics, 33(5), 896-902.

Li, X., & Sailor, D. J. (1995). Electricity use sensitivity to climate and climate change. World Resource Review, 7(3).

Luukkonen, R., Saikkonen, P., & Teräsvirta, T. (1988). Testing linearity against smooth transition autoregressive models. Biometrika, 75(3), 491-499.

Mideksa, T. K., & Kallbekken, S. (2010). The impact of climate change on the electricity market: A review. Energy Policy, 38(7), 3579-3585.

Moral-Carcedo, J., Vicéns-Otero, J. (2005). Modeling the non-linear response of Spanish electricity demand to temperature variations. Energy Economics, 27, 477–494.

Pardo, A., Meneu, V., Valor, E. (2002). Temperature and seasonality influences on Spanish electricity load. Energy Economics, 24, 55–70.

Peirson, J., & Henley, A. (1994). Electricity load and temperature: issues in dynamic specification. Energy Economics, 16(4), 235-243.

Petrick, S., Rehdanz, K., & Tol, R. S. (2010). The impact of temperature changes on residential energy consumption (No. 1618). Kiel working paper.

Quandt, R. E. (1988). The econometrics of disequilibrium. Blackwell.

Ranson, M., Morris, L., & Kats-Rubin, A. (2014). Climate Change and Space Heating Energy Demand: A Review of the Literature (No. 201407). National Center for Environmental Economics, US Environmental Protection Agency.

Sailor, D. J., Munõz, J. R. (1997). Sensitivity of electricity and natural gas consumption to climate in the USA — methodology and results for eight states. Energy, 22, 987–998.

Teräsvirta, T. (1994). Specification, estimation, and evaluation of smooth transition autoregressive models. Journal of the American Statistical Association, 89(425), 208-218.

Teräsvirta, T. (1996). Modelling economic relationships with smooth transition regressions (No. 131). Stockholm School of Economics.

Terasvirta, T., & Anderson, H. M. (1992). Characterizing nonlinearities in business cycles using smooth transition autoregressive models. Journal of Applied Econometrics, 7(S1), S119-S136.

Tsay, R. S. (1989). Testing and modelling threshold autoregressive processes. Journal of the American Statistical Association, 84(405), 231-240.

Tsay, R. S. (1998). Testing and modelling multivariate threshold models. Journal of the American Statistical Association, 93(443), 1188-1202.

Valor, E., Meneu, V., Caselles, V. (2001). Daily air temperature and electricity load in Spain. *Journal of Applied Meteorology* 408, 1413–1421.

Appendix A1: Scatter Plot of Filtered Consumption and Temperature (Filter I).

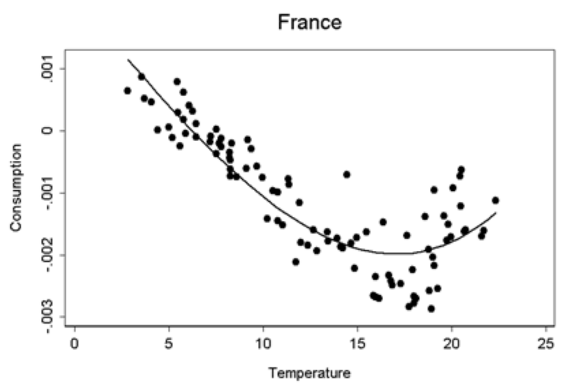

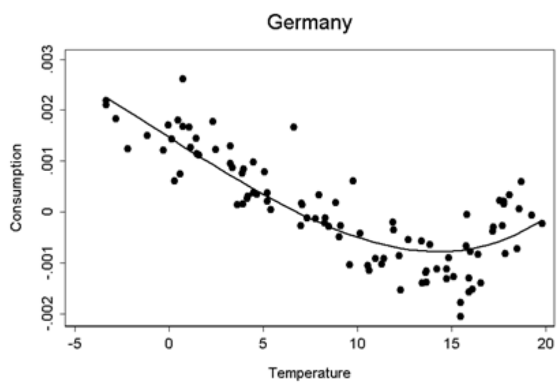

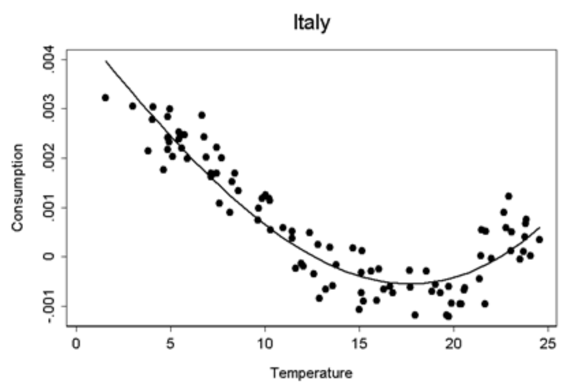

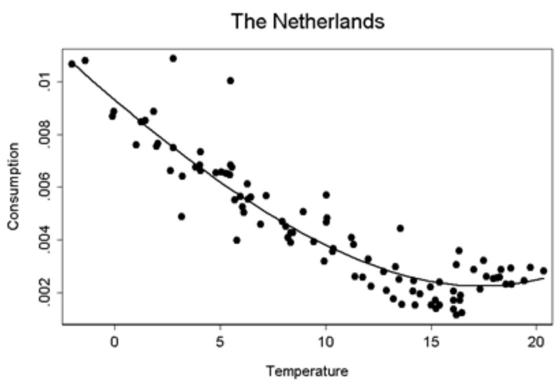

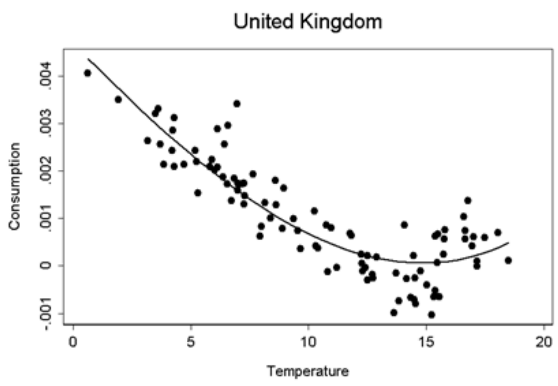

Notes: The fitted line portrays a local polynomial kernel regression of order 2 with a Gaussian kernel between the two variables.

Appendix A2: Scatter Plot of Filtered Consumption and Temperature (Filter II).

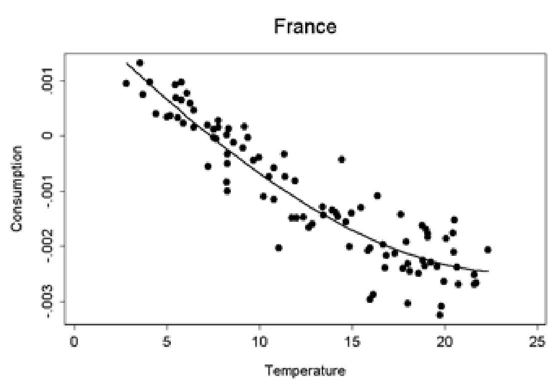

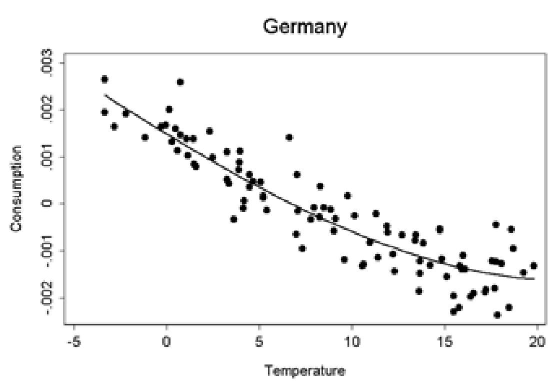

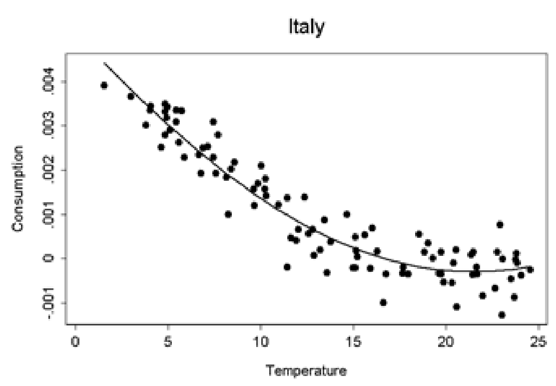

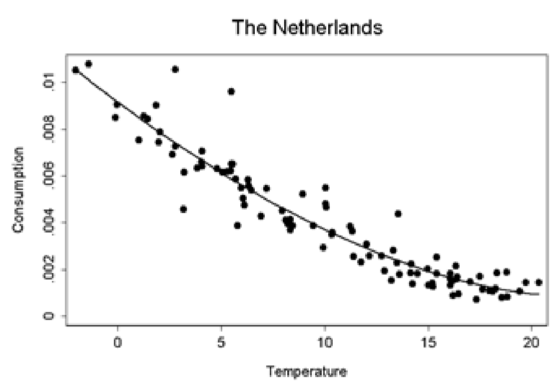

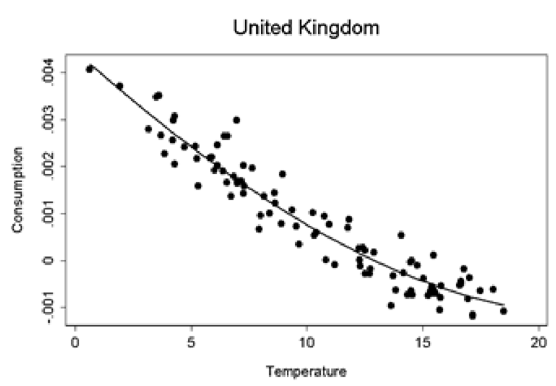

Notes: The fitted line portrays a local polynomial kernel regression of order 2 with a Gaussian kernel between the two variables.

Appendix A3: Scatter Plot of Regime Switching Probability and Temperature (Filter I).

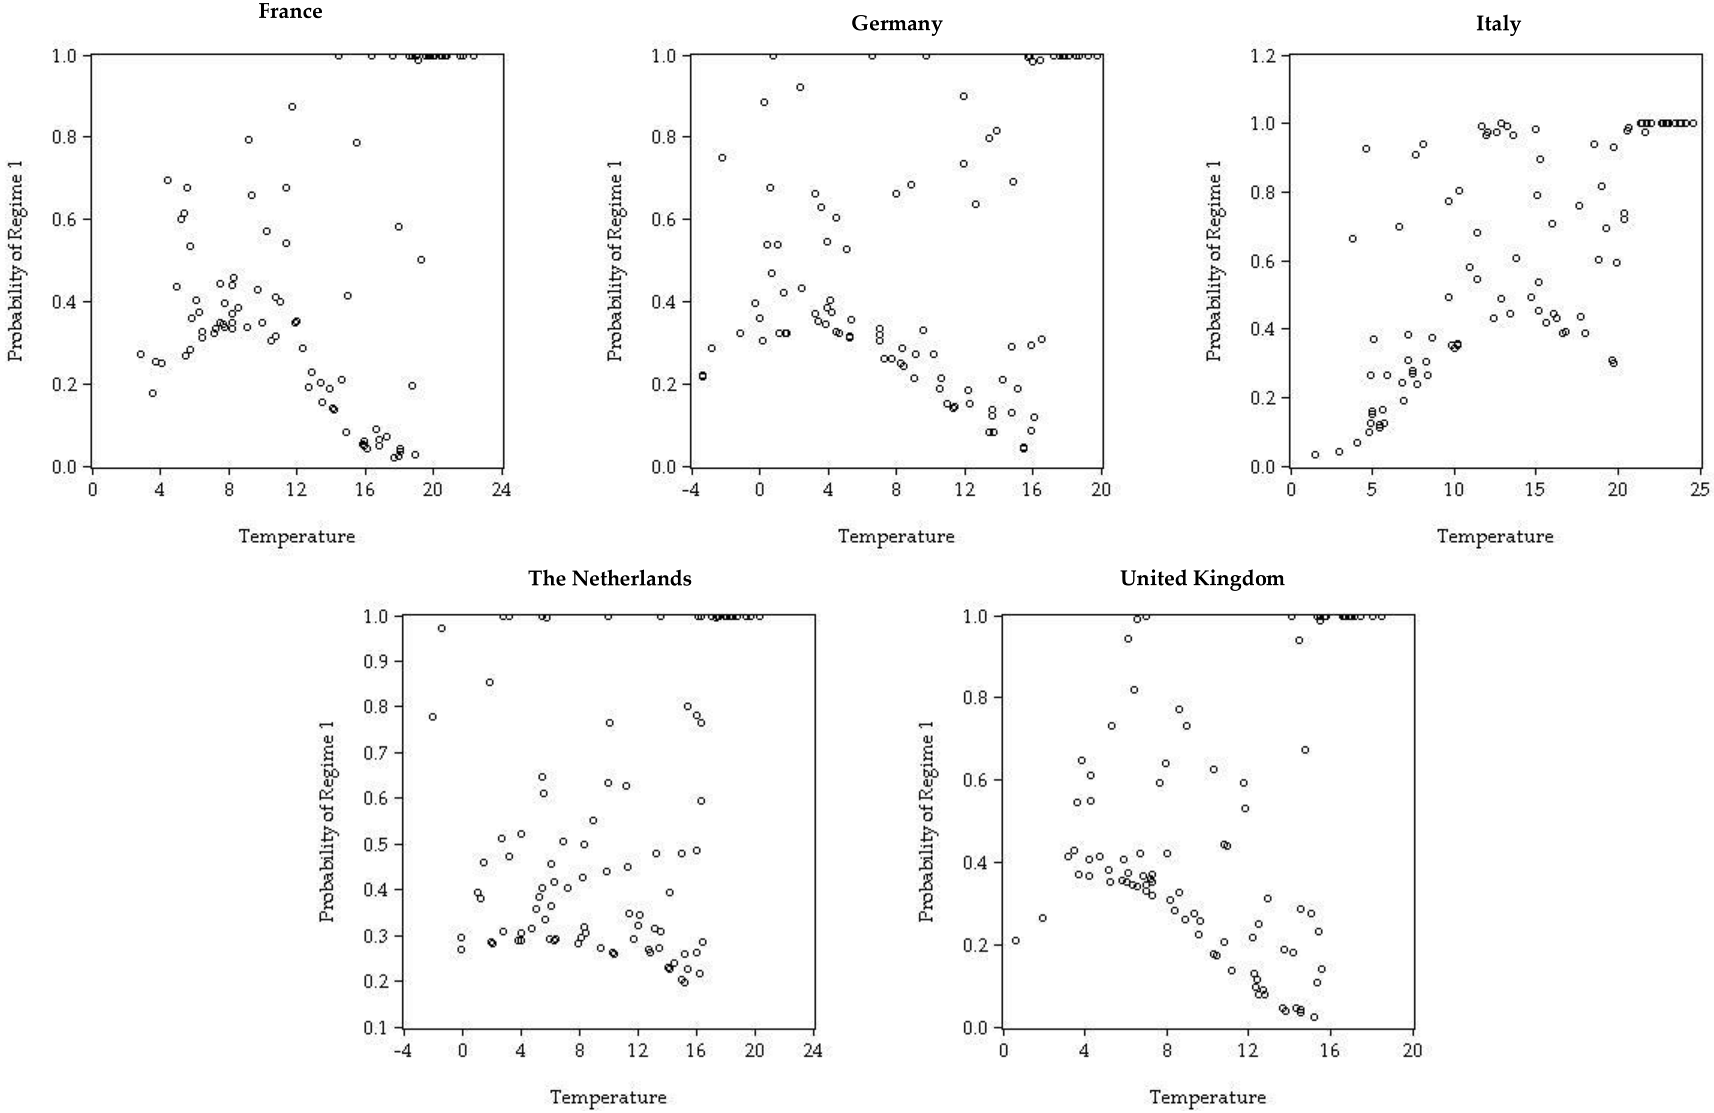

Notes: Portrays a scatter plot of the results of a simple regime-switching model with two states, without imposing any a priori relation between the switching probability and temperature.

Appendix A4: Scatter Plot of Regime Switching Probability and Temperature (Filter II).

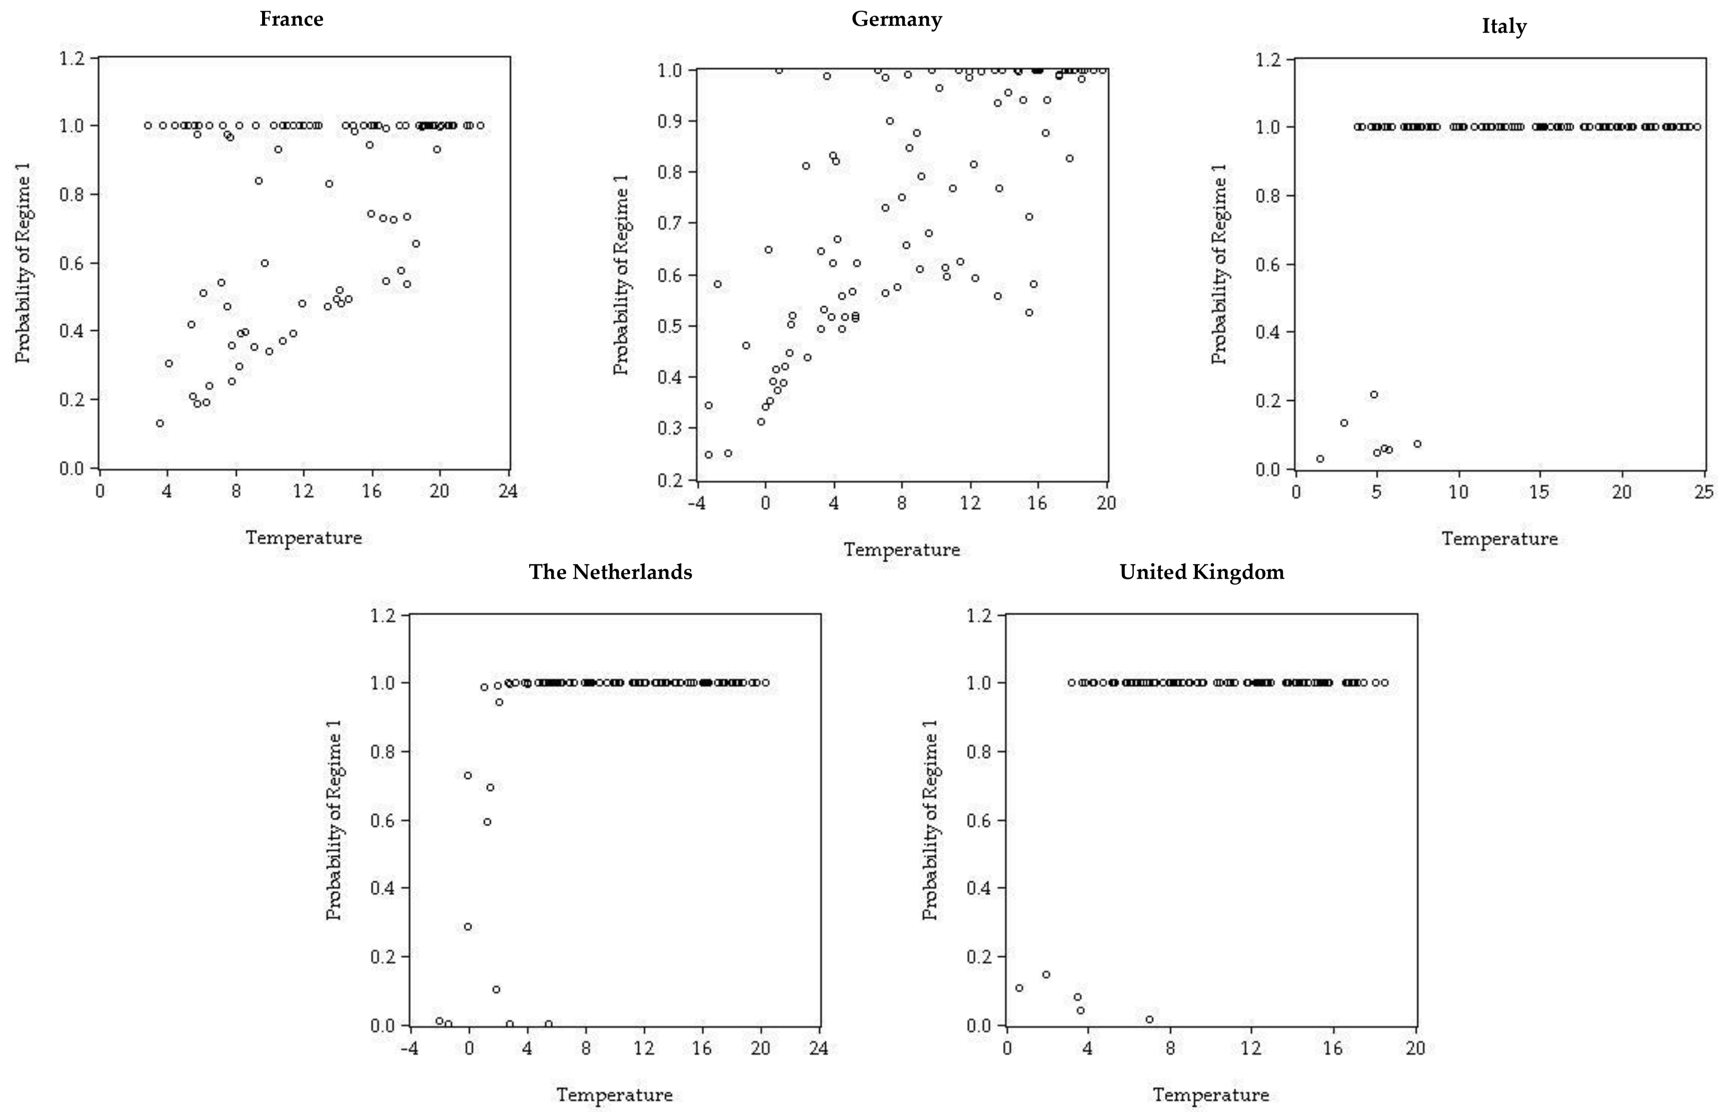

Notes: Portrays a scatter plot of the results of a simple regime-switching model with two states, without imposing any a priori relation between the switching probability and temperature.

Appendix A5: 3D-Plot of the AIC Criteria for the Two-Threshold Regression Model (France).

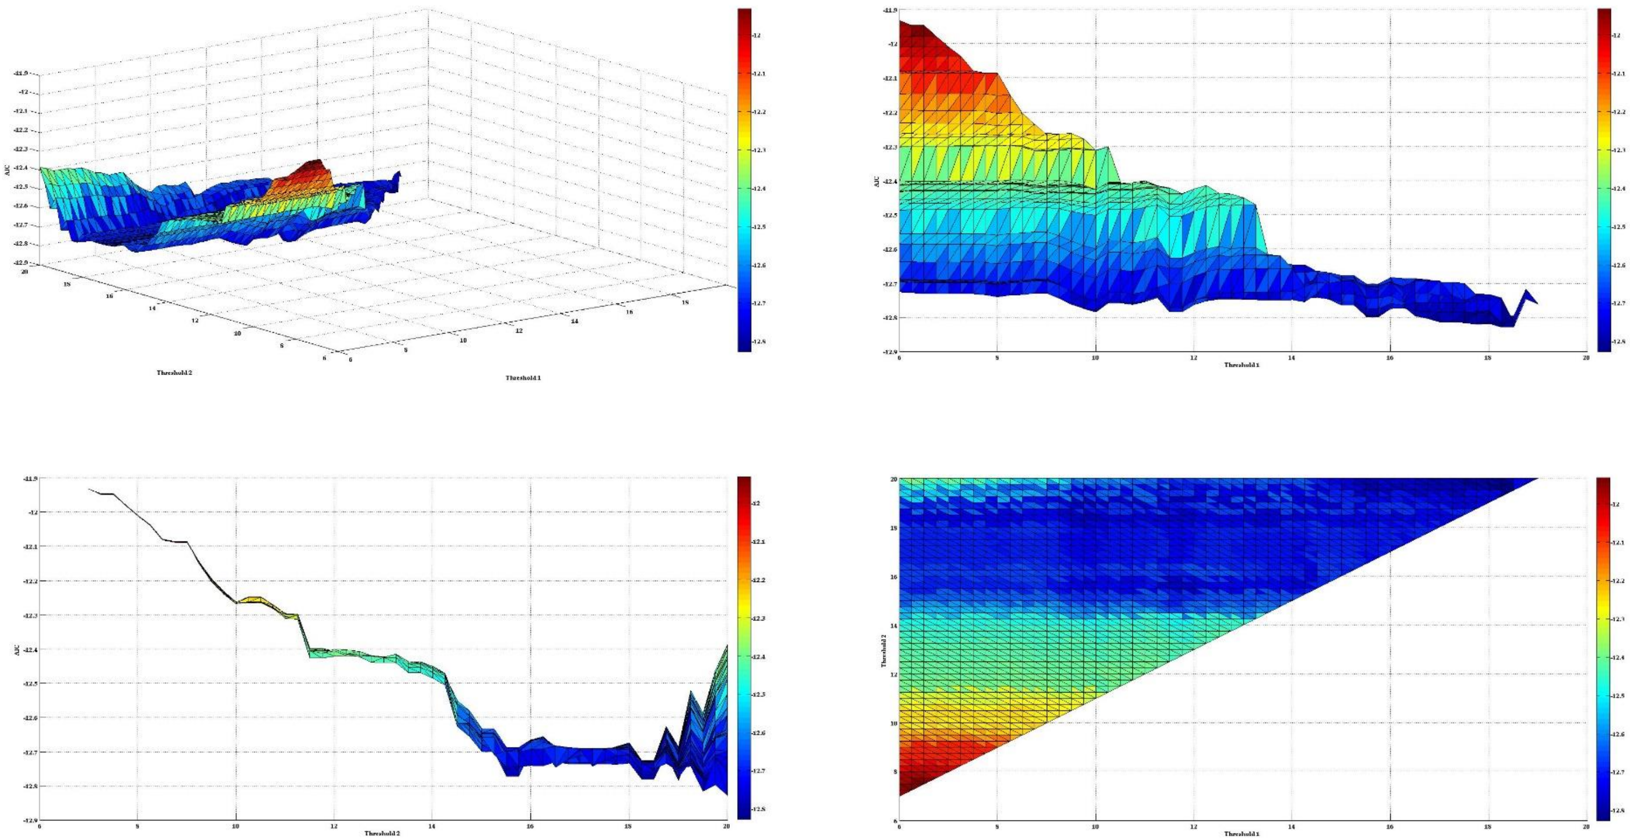

Notes: Depicts a 3D-plot of the results of a two-threshold regression model. In particular, a surface is estimated where the AIC serves as an evaluation criteria for different values of the two thresholds.

Appendix A6: 3D-Plot of the AIC Criteria for the Two-Threshold Regression Model (Germany).

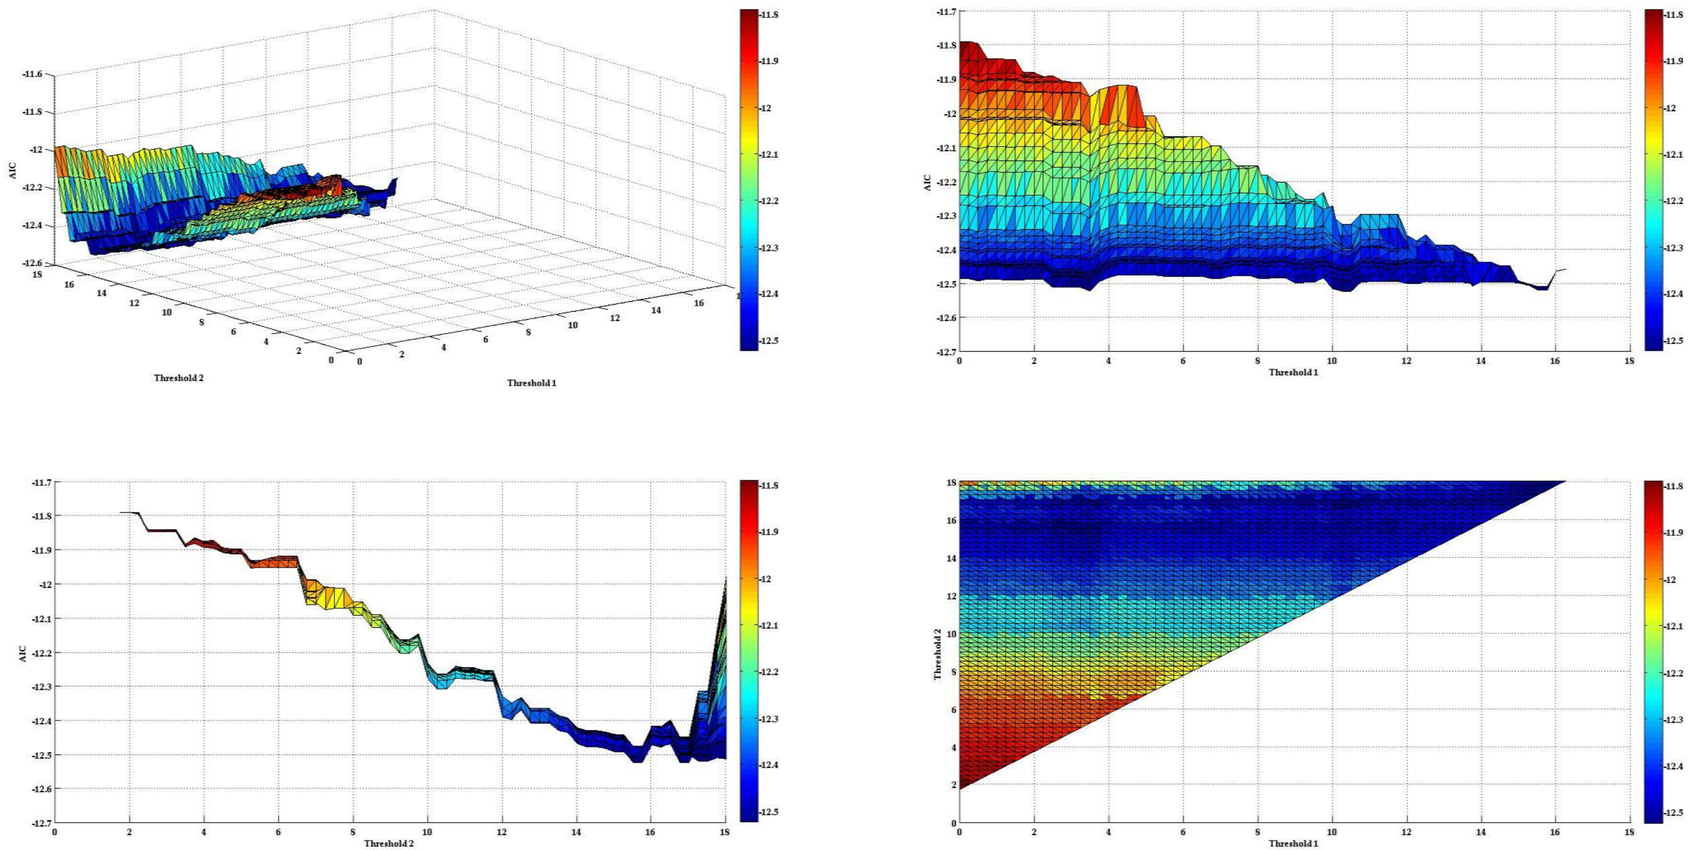

Notes: Depicts a 3D-plot of the results of a two-threshold regression model. In particular, a surface is estimated where the AIC serves as an evaluation criteria for different values of the two thresholds.

Appendix A7: 3D-Plot of the AIC Criteria for the Two-Threshold Regression Model (Italy).

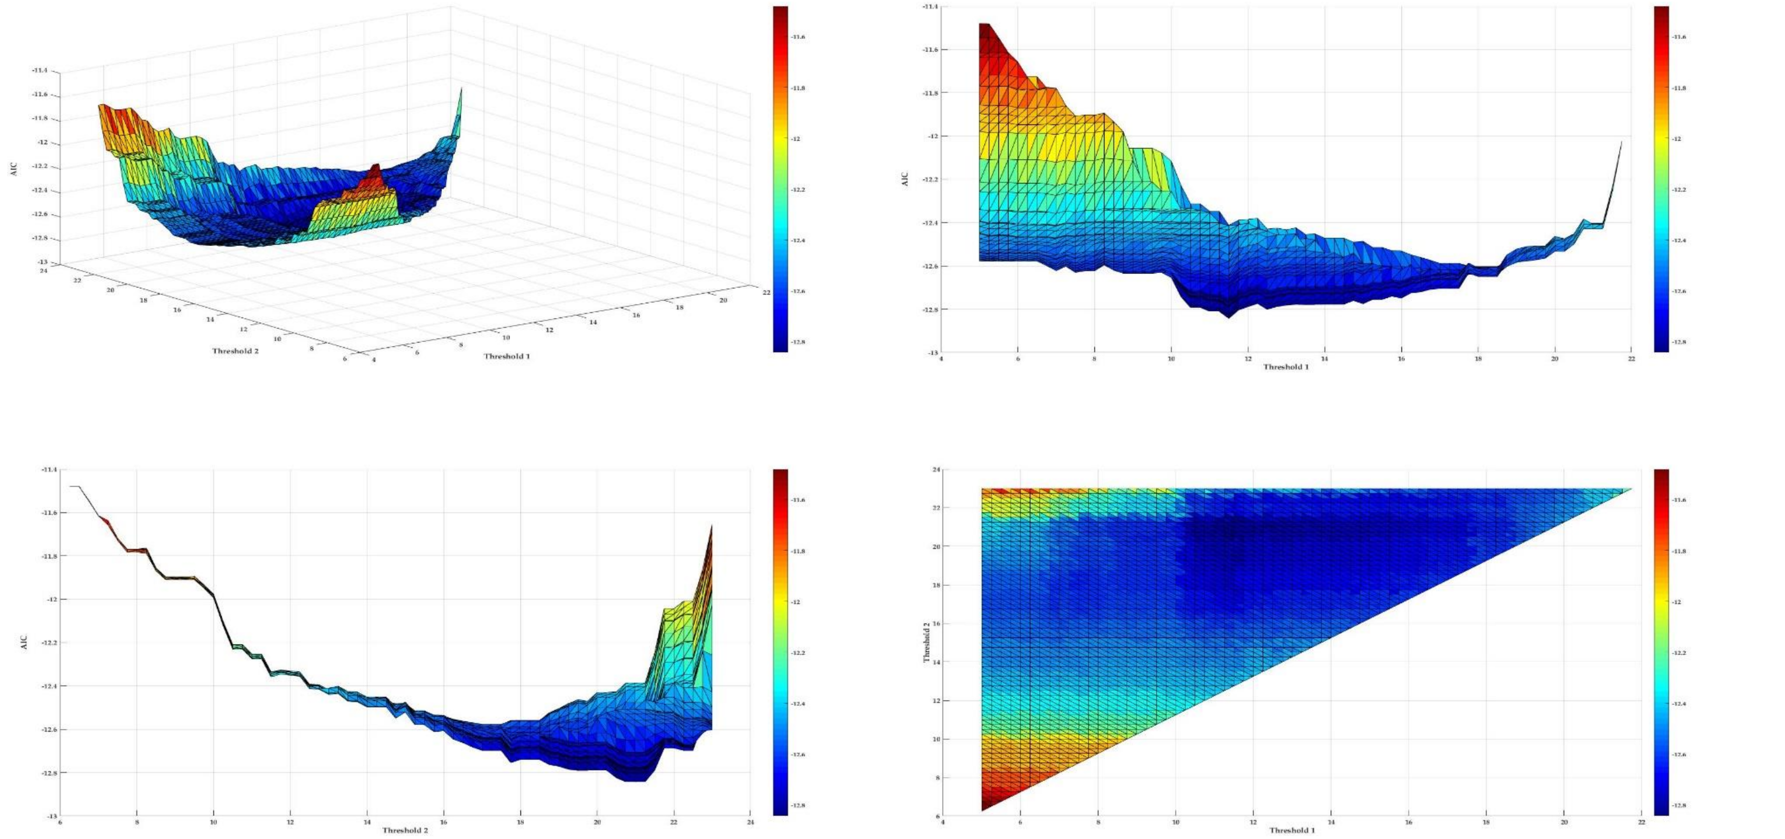

Notes: Depicts a 3D-plot of the results of a two-threshold regression model. In particular, a surface is estimated where the AIC serves as an evaluation criteria for different values of the two thresholds.

Appendix A8: 3D-Plot of the AIC Criteria for the Two-Threshold Regression Model (The Netherlands).

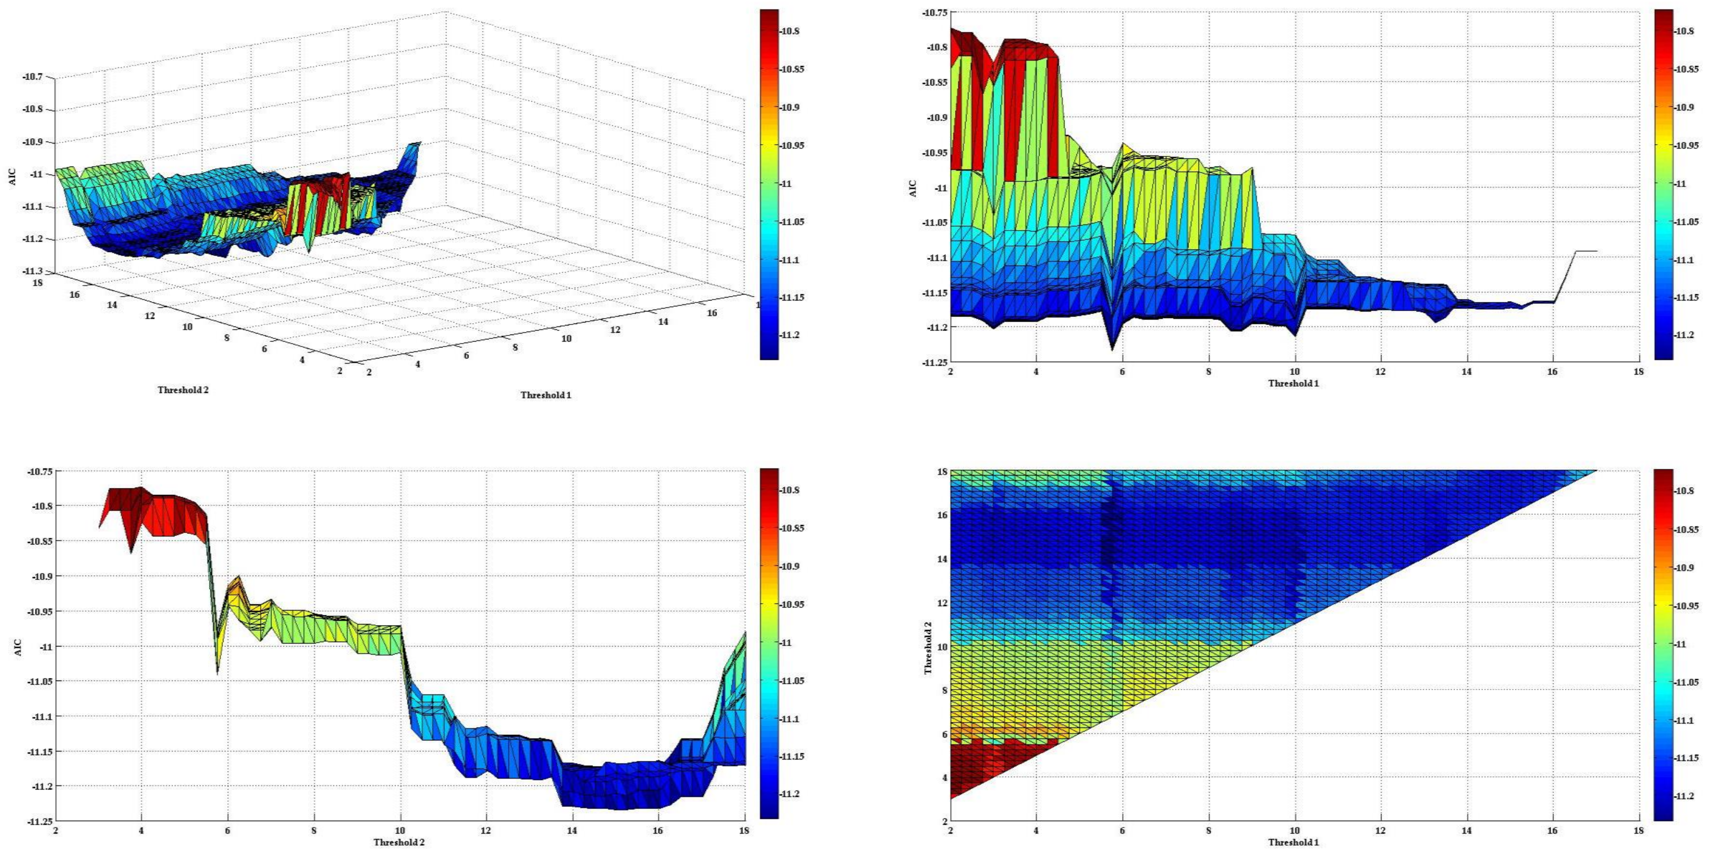

Notes: Depicts a 3D-plot of the results of a two-threshold regression model. In particular, a surface is estimated where the AIC serves as an evaluation criteria for different values of the two thresholds.

Appendix A9: 3D-Plot of the AIC Criteria for the Two-Threshold Regression Model (United Kingdom).

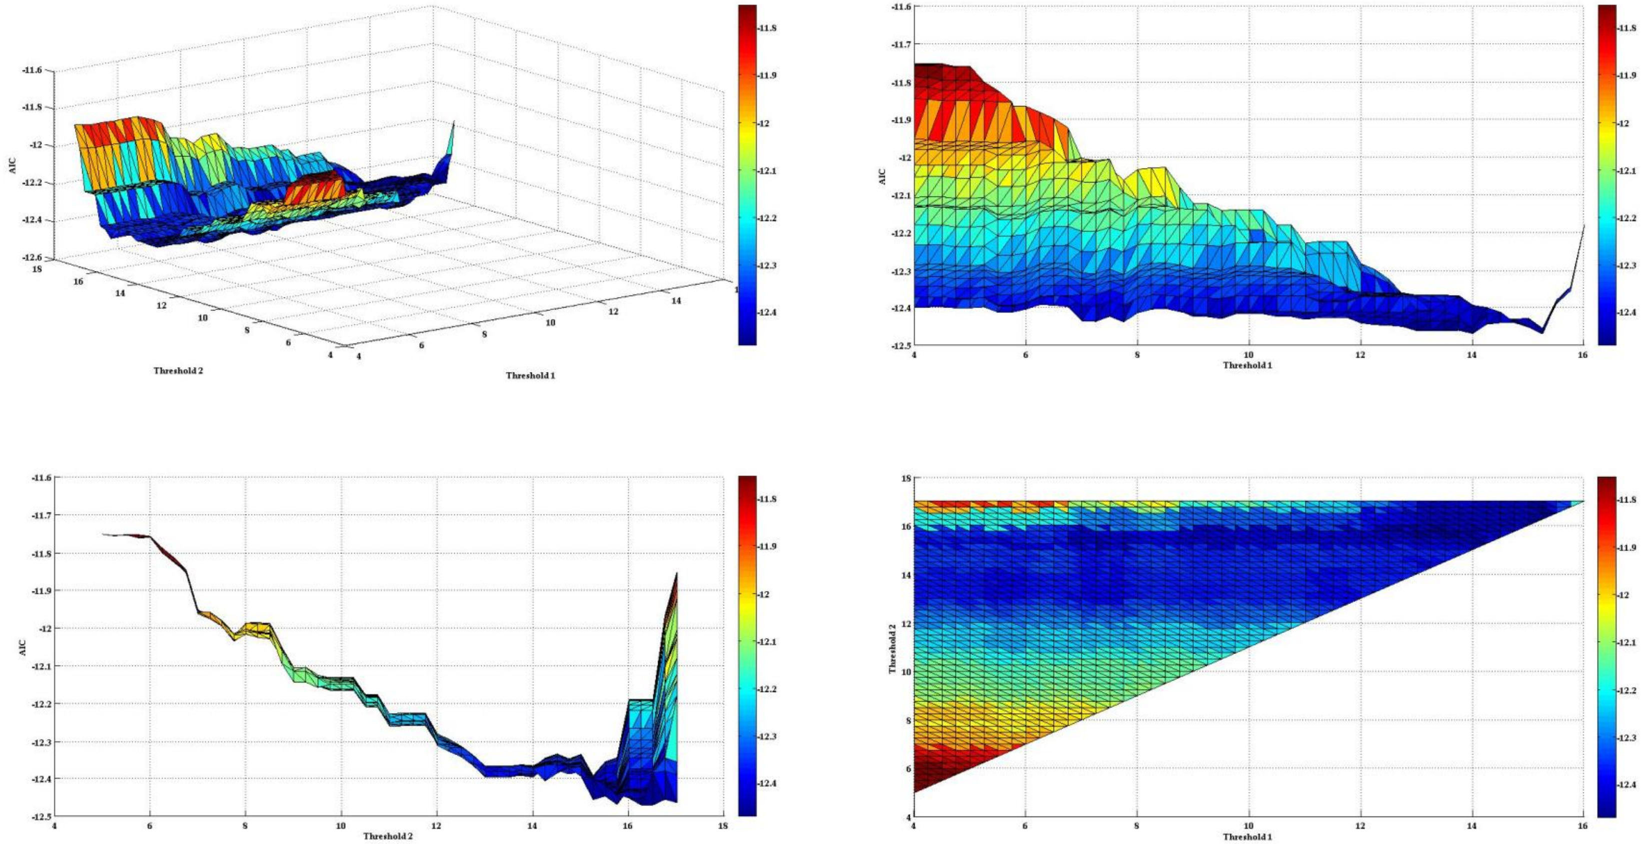

Notes: Depicts a 3D-plot of the results of a two-threshold regression model. In particular, a surface is estimated where the AIC serves as an evaluation criteria for different values of the two thresholds.