An Econometric Analysis of Electricity Demand Response to Price Changes at the Intra-Day Horizon: The Case of Manufacturing Industry in West Denmark.

Niels Framroze Møller[*] and Frits Møller Andersen

Technical University of Denmark, Management Engineering - Systems Analysis
August 18, 2015

ABSTRACT

The use of renewable energy implies a more variable supply of power. Market effciency may improve if demand can absorb some of this variability by being more flexible, e.g. by responding quickly to changes in the market price of power. To learn about this, in particular, whether demand responds already within the same day, we suggest an econometric model for hourly consumption-and price time series. This allows for multi-level seasonality and that information about day-ahead prices does not arrive every hour but every 24th hour (as a vector of 24 prices). We confront the model with data from the manufacturing industry of West Denmark (2007-2011). The results clearly suggest a lack of response. The policy implication is that relying exclusively on hourly price response by consumers for integrating volatile renewable electricity production is questionable. Either hourly price variation has to increase considerably or demand response technologies be installed.

1. Introduction

The purpose of this analysis is to gain insight into the dynamic response of electricity demand to price changes in the very short run. In particular, in this research we suggest a simple time series-based econometric approach to investigate whether hourly demand responds to hourly prices already within the same day. Based on this we analyse hourly time series of electricity consumption and prices for the manufacturing industry in West Denmark for the period 2007-2011. Our analysis is performed on two levels: we consider both the aggregate manufacturing industry as well as a single anonymous consumer. Most of the electricity consumers in the considered group have hourly metering and the option of hourly pricing.[1] Moreover, in the very short run – say hours – some industrial consumers are able to postpone electricity consumption without influencing their output significantly. E.g. a cold store may cut electricity consumption for an hour or two when prices are high. However, in order to maintain the temperature within the acceptable limit, electricity consumption must increase later. Other examples of industrial consumer flexibility relate to lighting, pumping and heating.

---

* Corresponding author e-mail: nfmo@dtu.dk

[1]   Although no information about the share of consumers using this option exists for the period under study, this is most likely to be large. In particular, it appears from an analysis made by the Danish Energy Agency in 2014 that, only around 25% of sales to the Danish industry come from fixed-price contracts (see www.ens.dk/info/nyheder/nyhedsarkiv/ny-metode-goer-virksomhedernes-elpriser-mere-retvisende and the links to background notes (in Danish) there).

Analysing short-term demand response is important for at least two reasons: First, if customers observe and react to hourly electricity prices, the efficiency of the electricity market is likely to improve and thus a welfare gain can be obtained [1, p. 70]. Second, production from renewable production technologies like wind and solar varies unsystematically and is only partly predictable. Technical integration of these technologies therefore requires continuous reactions either within other parts of the supply system or in the demand (e.g. demand response to hourly prices). In light of the future increases in the production share coming from the more volatile renewables, these arguments clearly become more and more significant.

The literature on short-term price response is extensive and estimated elasticities vary considerably reflecting both methodological differences and customer characteristics. Concerning load shifting, e.g. moving consumption from daytime to nights, time-of-use rates where customers know when rates shift from high to low, estimated elasticities are often both significant and relatively large, see e.g. [2] and [3]. Targeting periods with very high marginal production costs and market prices (either due to high demand or lack of production capacity), analyses of critical peak pricing schemes, where consumers are informed, typically a day in advance, that their rate will be extraordinary high in a specific period, also show significant and relatively large price elasticities. See e.g. [4] and [5].

For a successful integration of fluctuating renewable energy sources, customers need to react instantaneously to changing prices, typically hourly day-ahead market prices (plus taxes). Consumption should be increased when renewable production is large and prices low and decreased when renewable production is limited and prices high. Analyses of demand response to hourly prices show very mixed results with very large variations in the size of the price elasticity, but often reported price elasticities are fairly small and depend of specific customer characteristics [1, p. 88]. Estimated own-price elasticities vary from approximately zero to 0.38 (in absolute value) for a few large customers in peak periods.

The present research is based on standard econometric time series methods. However, the time series under study, in particular, the hourly Nord Pool prices, are a bit special, in that the information set of the agents is updated with new information on prices only every 24th hour and not every hour.[2] It can be shown that, in general this has to be taken into account for the estimation results to be reliable. We suggest a simple solution to this problem which implies a rearrangement of the original time series of consecutive hourly observations. The basic idea, which builds on [6], is to divide the day into a number of sub-periods, for example (but not necessarily) the 24 hours. In this way each new observation should rather be viewed as a vector or a panel of 24 variables, namely the 24 sub-period price and consumption levels. Thus, each new observation corresponds to a new day (rather than a new hour), i.e. when the information set updates with respect to prices.

To analyse intra-day price responsiveness based on the rearranged time series we suggest a simple, albeit general, "structural" or behavioural framework from which we derive a regression model for each sub-period's consumption level. The latter is regressed on prices from *all* sub-periods. In our empirical application, for example, we divide the day into 12 two-hour sub-periods which implies that we estimate twelve separate regression equations. We assume that agents are price takers, in that, electricity prices are determined by aggregate demand and supply which are approximately uninfluenced by the consumption unit we look at.

As is well known hourly electricity data display a rather pronounced degree of multi-level seasonality, i.e. periodic systematic patterns over the day, week and year, [7, Chapter 2]. For simple regressions, involving only the levels of consumption and prices, it is inevitable that a large part of this seasonality remains in the error term. To accommodate this, the regression model we use is therefore allowed to have multiplicative seasonal ARIMA errors. This model is denoted as the RegARIMA, [8]. In addition to the seasonal dynamics, this model also allows for non-seasonal dynamics, i.e. the usual AR and MA terms. Such terms are also required to capture the high degree of inter-day correlations for both consumption and prices.

Taking our model to the data, the results clearly suggest a lack of demand response to price changes at the intra-day horizon. This holds for both the aggregate

---

[2] In general, the *information set at time t*, say Ω_t, has as elements the variables that are known to the agent at time *t*. Econometrically these are the variables we condition on.

manufacturing sector of West Denmark as well as for a single anonymous consumer from this industry. This conclusion is obtained in statistically well-specified RegARIMA models and is obtained independently of whether these are identified manually or by automatic ARIMA modeling algorithms.[3, 4]

In the next section we first elaborate briefy on the above-mentioned temporal aspect of the price series. Then we introduce the notation and technical details, present a structural framework and from all this, derive the RegARIMA model to be estimated. We confront our RegARIMA model with the data in Section 3 and finally conclude and discuss our findings in Section 4.

2. Modelling hourly electricity demand based on dynamic time series models

Assume that the data at hand come in the form of time series data of electricity prices and load with an hourly resolution. If one is to apply standard statistical models to these time series, in order to analyse electricity demand as a function of electricity prices, there are a few fundamental characteristics to be taken into account.

First, although electricity is priced on an hourly basis, the information set (with respect to prices) of the consumption unit is not updated each hour. Instead, the 24 hourly prices, corresponding to electricity delivery for each of the 24 hours on a given day, have been determined simultaneously in an auction taking place the day before. That is, prices are determined in a day-ahead market, in this case, the Nord Pool market.[5] This is in contrast to the assumption underlying many applied time series analyses, namely that the information set is updated with new information each new period. This characteristic temporal property of day-ahead electricity prices has been emphasized before in connection with analyses of prices only, i.e. not jointly with electricity consumption, as is the case here, and different ways to deal with it have been suggested.[6] For example, panel models have been suggested for which the price series is treated as a panel of 24 cross sectional units corresponding to the 24 hourly prices measured each day [9]. Another way to handle day-ahead price series, which has been suggested by Wolak, is simply to stack the 24 prices into a vector and then treat this vector as a time series process [6]. The present analysis builds on the latter idea and generalises it to apply to a structural model involving consumption and prices (and possibly other variables).

In both of the above-mentioned studies the way to take the temporal property into account is to transform or rearrange the original time series so that it is amenable to standard time series models. This is also the approach here and within our approach an observation of prices is basically a vector of say 24 (hourly) prices. The same goes for consumption of course. For example, if we consider 3 days (hour 1, day 1 to hour 24 day 3), that is, 72 hours, we have a multivariate series with three 24-dimensional observations, instead of a univariate series with 72 observations. Alternatively, if we divide the day into 12 sub-periods we would still have three observations but these would be 12-dimensional. In this setup intra-day effects from prices to consumption are thus formally treated as static or current effects.

Although these intra-day effects are of primary interest here the model must also allow for dynamic dependence, i.e. inter-day dependence, since consumers may be able to shift their consumption across days. Moreover, expectations of prices are likely to be adaptive and thus related to past prices, and there may be physical restrictions which imply, for example, that reduced electricity consumption on a given day means increased consumption the following day. These aspects mean that the model must also allow for dynamic dependence beyond a day's length.

In addition to this time dependence there is also a strong multi-level seasonality in the original hourly series, that is periodicity over the day, the week and the year (see e.g. [7]). Note that, in our transformed series (into multivariate daily series) only the weekly and yearly seasonality remain.

In order to take account of all this, i.e. the special temporal property of prices series and both the seasonal and non-seasonal dynamics, we analyse the "sub-period transformed" time series by use of a linear regression model with a multiplicative seasonal ARIMA error term. The latter regression model is often referred to as the RegARIMA model [8]. The RegARIMA is relatively general as it comprises a range of time series

---

[3]   By *well-specified* we mean that the most important residual-based diagnostics or misspecification tests were passed to a satisfactory extent.
[4]   See [17].
[5]   See http://www.nordpoolspot.com/.
[6]   See e.g. [11] for a clear discussion and for further references.

models used in the literature on short-term load forecasting. These include the linear regression model with white noise errors, pure AR models, pure MA, regression with MA-or AR errors, pure seasonal MA and AR models, and of course various combinations, such as a linear regression with seasonal autoregressive errors (see e.g. [7, Section 3.4], for a survey).

To a large extent our approach follows what has been a tendency in the literature on short-term load modelling and forecasting since [10], namely to have a distinct model for each hour (or sub-period) of the day. In particular, note that, although our analysis is akin to that of [10], these authors do not include prices as regressors but are concerned with load only. But clearly setting $I = 24$ and choosing a suitable ARIMA structure this will reproduce their model. However, the RegARIMA algorithm that we apply here allows us to estimate a richer dynamic structure for the error process.

In the next section we first provide some technical details about the basic transformation underlying the econometric model. This serves the purpose of both introducing the notation but should also facilitate the exposition and make our analysis more transparent and thus easier for other researchers to apply. Then we provide a general behavioural foundation for the econometric model (to facilitate the interpretation of the empirical results later on), and from this derive the RegARIMA regression for sub-period consumption.

2.1. The econometric approach
2.1.1. The time axis, measurements and notation
We consider a time axis where the unit of measurement is one day indexed by t, and divide day t into I non-overlapping *sub-periods*, $I \in \{1, 2, 3, ..., 24\}$. Day $t = 1$, 2, 3, …, do not necessarily have to be consecutive calender days, and may for example, exclude weekends and holidays. However, in the application below all calender days are analysed consecutively. The sub-periods are not necessarily of the same length but their length is always an integer number of hours. For example, the first period could be the first hour and the second period could be the remaining hours of the day. But of course there are many other possibilities. We order the periods corresponding to $i = 1, 2, .., I$, chronologically, but it is not required that these periods are adjacent. The sum of these periods can be at most 24

hours but may clearly be less. It thus follows that, if $I = 24$, period $i = 1, 2, .., I$ corresponds to the 24 consecutive hours of the day, and that, $I > 1$ is required, if we are to allow for intra-day effects which is the focus here. Note, for example, that $I = 1$ could correspond to a whole day (24-hour period) or even a single given hour of the day.

To keep the exposition simple we assume in the remainder of this section that $I = 2$, *for example* dividing the 24 hours into two 12-hour periods. In terms of the illustration here, there is no loss of generality in making this simplifying assumption. Of course, in a given empirical application it may be preferable not to "aggregate too much over time" by letting each of the I periods correspond to several hours, since this is likely to hide potentially interesting dynamic effects. In the empirical analysis in Section 3 we therefore allow for as many as 12 sub-periods ($I = 12$), where each sub-period corresponds to a two-hour period.

Now, given hourly observations on consumption and prices of electricity for sub-period i of day t, we compute the variables, $C_{i,t}$ and $P_{i,t}$. In general, these variables are thus functions of the original hourly series. For example, if sub-period i consists of three hours, $C_{i,t}$ could be the aggregate or average consumption for these three hours, and $P_{i,t}$ could be the average price. In the application in Section 3, $C_{i,t}$ denotes the aggregate consumption over the hours corresponding to sub-period i, while $P_{i,t}$ is the average price for this period. In this paper, we let capital letters denote the original variables to be distinguished from the logarithmic values which we denote by small letters. The reason for the logarithmic transformation is that the regression models in Section 3 can be viewed as log-linear approximations to more general non-log-linear equations. Moreover, by use of a range-mean plot, based on sub-samples of the seasonal length = 7, we found that a logarithmic transformation of the variables was in fact clearly supported by the data.

Already now it should, to some extent, appear that there is some generality in our approach, in that there are a large number of possible (and interesting) combinations of the number and length of sub-periods, that one could experiment with. Note also that, studies of the daily time series of consumption in a particular hour [11], or an average computed for the day, are examples implying $I = 1$.

2.1.2. A structural econometric model for sub-period demand and the RegARIMA

To facilitate the interpretation of the empirical results we need a structural or behavioural model. We want it to represent different electricity consumption units and thus suggest here a relatively general formulation. Its substantive structure is however straightforward: Each period t (day), after the information set updates with respect to prices, consumption is planned for the current period, i.e. all the sub-periods of period t, and possibly a number of future sub-periods. Given the past (lagged values of the relevant variables), planning is based on the current realized prices, which are treated as exogenous by the price taking assumption, and supposedly other exogenous variables, such as planned output and (expected) temperature etc. In addition, expected future values of prices and of the exogenous variables also influence the planning.[7]

As a simple illustration assume that planning only goes as far as one day ahead. That is, simultaneous planning for period t and $t + 1$ or rather the four sub-periods, 1 and 2 of day t and 1 and 2 of day $t +1$. If we let the superscripts, p and e, refer respectively to, planned and expected values, to be distinguished from realized values, the planning equations can be written as,

$$
\begin{aligned}
C_{1,t}^p &= f_1\left(C_{2,t}^p, C_{1,t+1}^p, C_{2,t+1}^p, P_{1,t+1}^e, P_{2,t+1}^e, C_{1,t-1}, C_{2,t-1}, P_{1,t}, P_{2,t}, \mathbf{X}_t^{'}\right), \\
C_{2,t}^p &= f_2\left(C_{1,t}^p, C_{1,t+1}^p, C_{2,t+1}^p, P_{1,t+1}^e, P_{2,t+1}^e, C_{1,t-1}, C_{2,t-1}, P_{1,t}, P_{2,t}, \mathbf{X}_t^{'}\right), \\
C_{1,t+1}^p &= f_3\left(C_{1,t}^p, C_{2,t}^p, C_{2,t+1}^p, P_{1,t+1}^e, P_{2,t+1}^e, C_{1,t-1}, C_{2,t-1}, P_{1,t}, P_{2,t}, \mathbf{X}_t^{'}\right), \\
C_{2,t+1}^p &= f_4\left(C_{1,t}^p, C_{2,t}^p, C_{1,t+1}^p, P_{1,t+1}^e, P_{2,t+1}^e, C_{1,t-1}, C_{2,t-1}, P_{1,t}, P_{2,t}, \mathbf{X}_t^{'}\right).
\end{aligned}
\tag{1}
$$

where the f functions are differentiable. Here, $\mathbf{X}_t$ is a vector of exogenous variables (other than prices), i.e. those that are given in the planning problem. Expected values of such variables as well as indictors for the type of the day (work day, holiday etc.), sinusoidal functions capturing annually seasonality etc. can also be included in $\mathbf{X}_t$ (see Section 3). We may think of this system of equations as sufficient first order conditions corresponding to some underlying representative optimization problem (e.g. cost minimization).

There are a few additional aspects of Eq. (1) to note. First, it is reasonable that lagged terms, $C_{1,t-1}$ and $C_{2,t-1}$, enter. This could for example reflect physical/technical constraints (fridge/cooling heating) and/or expected load requirements: What is used in the most recent period is likely to influence consumption in the current period. Secondly, note that for both past consumption and lagged exogenous variables (included in $\mathbf{X}_t$), additional lags are likely needed to obtain an empirically viable model. Third, there is no direct effect on the plans from the lagged prices. Clearly, as shown below, there may indeed be an indirect effect, working through the (adaptive) price expectations. Fourth, note that, here the horizon for expectations is two sub-periods and hence coincide with the planning horizon, which is not generally the case. Fifth, it is also an arbitrary assumption (made for the sake of illustration) that the consumption unit is planning only one day ahead (i.e. planning simultaneously for period t and $t + 1$). The planning horizon could be longer and, moreover, it could imply, say only half of day $t + 1$, for example.

To close the model we assume that all expectations are adaptive. Specifically, since prices for different sub-period of day t are correlated and prices may be correlated over days, we allow the expected price level for a given sub-period to depend on the prices in several sub-periods. For example, $P_{1,t+1}^e$ may depend on $P_{1,\ t}$, $P_{2,t}$, $P_{1,t-1}$, $P_{2,t-1}$ and $P_{1,t-6}$ etc. Moreover, current and lagged values of the other variables may also influence price expectations: The important assumption is that all expectations are functions of only current and lagged values of the observable variables. We assume furthermore that these functions are differentiable.

To come from the planning equations in Eq. (1) to the optimal planned levels, equations for the expected values, which fulfil adaptiveness, are inserted in Eq. (1). Assuming, by regularity, that the Jacobian matrix of first-order derivatives is non-singular, the Implicit Function Theorem ensures that the resulting system of equations can be solved with respect to $C_{1\ ,t}^p, C_{2,t}^p, C_{1\ ,t+1}^p$ and $C_{2\ ,t+1}^p$ which gives the optimally planned values. See e.g. [12].

To come from these optimal planned magnitudes to the equations for the observable variables, which form the basis for the estimation equations below, we could assume the following observation mappings,

$$
C_{1,t} = g_1\left(C_{1,t}^p, V_{1,t}\right)
$$

$$
C_{2,t} = g_2\left(C_{2,t}^p, V_{2,t}\right)
\tag{2}
$$

where the g functions are differentiable and the $V_{.,t}$ terms in Eq. (2) are unsystematic, i.e. white noise, error

---

[7]   Thus, using the terminology of [18], a *plan* is here a mix of a contingent plan and a behavioral model (i.e. based on expected values). See [18, Chapter 6].

components. We assume that the planned values, $C_{1,t+1}^p$ and $C_{2,t+1}^p$ in Eq. (1), are not binding and are not used. They are discarded since the information set is updated and new plans can be made at the beginning of period $t+1$.

Instead of Eq. (2) we shall allow for slightly more flexibility of the specification, by assuming realistically that the consumer is free to adjust his sub-period consumption when entering that sub-period. For example, as seen from the first equation in Eq. (2), $C_{1,t}$ deviates from $C_{1,t}^p$ due to the error term representing unforeseen events. Since consumption $C_{1,t}$ is most likely influencing $C_{2,t}$ the deviation of $C_{2,t}$ from the planned level, $C_{2,t}^p$, is likely to depend also on the realised value, $C_{1,t}$ and not only $V_{2,t}$. This leads us to the alternative recursive formulation,

$$C_{1,t} = g_1(C_{1,t}^p, V_{1,t}),$$
$$C_{2,t} = g_2(C_{2,t}^p, C_{1,t}, V_{2,t}), \tag{3}$$

for which it is assumed that there are no effects of $C_{1,t}$ on $C_{2,t}$ when $C_{1,t} = C_{1,t}^p$.

To derive the estimable regression equations, note first that it follows from the assumption of adaptive expectations that the solutions of Eq. (1), $C_{1,t}^p$ and $C_{2,t}^p$ in Eq. (3), depend only on lagged consumption, current and lagged values of prices and other exogenous variables. Assuming a log-linear form of the g functions in Eq. (3), or more generally, making a log-linear approximation, we take logs in Eq.(3),

$$c_{1,t} = x_{1,t} + \beta_{1,1}p_{1,t} + \beta_{1,2}p_{2,t} + \beta_{1,3}p_{1,t-1} +$$
$$\beta_{1,4}p_{2,t-1} + \beta_{1,5}c_{2,t-1} + u_{1,t}$$
$$c_{2,t} = x_{2,t} + \beta_{2,1}p_{1,t} + \beta_{2,2}p_{2,t} + \beta_{2,3}p_{1,t-1} + \tag{4}$$
$$\beta_{2,4}p_{2,t-1} + \beta_{2,5}c_{2,t-1} + u_{2,t}$$

where $x_{\cdot,t}$ comprise the observable variables in $\boldsymbol{X}_t$, including lagged variables as a result of adaptive expectations, and additional lagged observable variables, e.g. higher lags of consumption. The reason why we stress "observable" here is that, in practice data on many of the exogenous variables are not available/observable at an hourly resolution. Hence, in the empirical application the influence from these variables is hidden in the error term. The $u_{1,t}$ and $u_{2,t}$ error terms are therefore allowed to be systematic, in

particular to follow an ARIMA structure describing both non-seasonal and seasonal dynamics (see below).

In the general case, for which I is no longer restricted to 2, we may state the regression equation for sub-period i consumption, $c_{i,t}$, in terms of the RegARIMA (linear regression with ARIMA error structure) formulation,

$$c_{i,t} = \boldsymbol{\beta}_i' \boldsymbol{Z}_{i,t} + u_{i,t}, \tag{5}$$

for $t = 1, 2, .., T$ and $i = 1, 2, ..., I$ and where, we have stacked all indicators and deterministic terms, current prices and lagged prices, lagged consumption and exogenous variables into the vector $\boldsymbol{Z}_{i,t}$. The equation for the ARIMA error structure is,

$$\Theta_{i,p_i}(L)\Psi_{i,P_i}(L^s)\Delta^{d_i}\Delta_s^{D_i}u_{i,t} = \Gamma_{i,q_i}$$
$$(L)\Pi_{i,Q_i}(L^s)\varepsilon_{i,t,} \tag{6}$$

for $t = 1, 2, .., T$ and $i = 1, 2, .., I$, and where L is the lag operator, $\Delta^{d_i} \equiv (1 - L)^{d_i}$, $\Delta_s^{D_i} \equiv (1 - L^s)D^i$ with d_i and D_i being integers and $s = 7$ corresponding to the weekly seasonality. We assume that $\varepsilon_{i,t}$ and $\varepsilon_{j,s}$ are uncorrelated for $t \neq s$ for all i and j, and for $i \neq j$ when $t = s$. This implies that although all the I regression equations constitute a system one may still rely on single-equation estimation (i.e. equation by equation). See e.g. [13, Chapter 12]. Note that the uncorrelatedness between $\varepsilon_{i,t}$ and $\varepsilon_{j,t}$ for $i \neq j$ results since the system has a recursive structure, in that $C_{1,t}$ does not depend on $C_{i,t}$ for $i > 1$, but $C_{2,t}$ depends on $C_{1,t}$, $C_{3,t}$ depends on $C_{1,t}$ and $C_{2,t}$, $C_{4,t}$ depends on $C_{1,t}$, $C_{2,t}$ and $C_{3,t}$ and so forth. This recursive structure is due to the fact that sub-period i precedes sub-period $i + 1$, which precedes sub-period $i + 2$ etc. and the assumption that the consumer is able to adjust his sub-period consumption when entering that sub-period, i.e. the "ex post plans" adjustment as mentioned in connection with Eq. (3).[8]

The various lag-polynomials can be divided into those describing the *non-seasonal dynamics*,

$$\Theta_{i,p_i}(L) \equiv (1 - \theta_{i,1}L - \theta_{i,2}L^2 - ... - \theta_{i,p_i}L^{p_i}),$$
$$\Gamma_{i,p_i}(L) \equiv (1 - \gamma_{i,1}L - \gamma_{i,2}L^2 - ... - \gamma_{i,q_i}L^{q_i}),$$

and into those describing the *seasonal dynamics*,

$$\Psi_{i,p_i}(L^s) \equiv (1 - \psi_{i,1}L^s - \psi_{i,2}L^{s2}... - \psi_{i,p_i}L^{sP_i}),$$
$$\Pi_{i,Q_i}(L^s) \equiv (1 - \pi_{i,1}L^s - \pi_{i,2}L^{s2}... - \pi_{i,Q_i}L^{sQ_i}).$$

---

[8]   One may add that, at least to our knowledge, a fully developed software (including diagnostics test etc.) for estimating the corresponding RegARIMA *system*, that is, a multiplicative VARIMA with exogenous regressors, is not available.

Note that when an order index (p_i, P_i, q_i, or Q_i) is zero the corresponding polynomial is equal to 1, and that in practice (including the present analysis) these indices are often found to be relatively low, so that a relatively flexible dynamic structure can be described by a few parameters.

All four lag polynomials have (non-explosive) roots whose moduli are located in the complex plane such that the differenced process is stationary and invertible see e.g. [14]. The term, Δ^{d_i}, corresponds to d_i real unit roots, i.e. located in (1,0) in the complex plane. These roots imply non-stationarity of the integrated type, i.e. that can be removed by (first-) differencing (see e.g. [15]). The term $\Delta_s^{D_i}$ corresponds to non-stationarity in the form of $D_i \times s$ unit roots, which are spread out evenly on the unit circle (see e.g. [16]). These are referred to as *seasonal unit roots* and they represent non-stationarity that can be removed by taking the seasonal difference D_i times.

The RegARIMA model can also be stated by inserting Eq. (5) in Eq. (6), i.e.

$$\Theta_{i,p_i}(L)\Psi_{i,Pi}(L^s)\Delta^{d_i}\Delta_s^{D_i}\left(c_{i,t} - \beta_i'\right)$$

$$\mathbf{Z}_{i,t} = \Gamma_{i,q_i}(L)\Pi_{i,Q_i}(L^s)\varepsilon_{i,t},$$

and is symbolically written as the *multiplicative* seasonal RegARIMA denoted, $RegARIMA(p_i, d_i, q_i) \times (P_i, D_i, Q_i)_s$, resembling standard notation. Finally, yet another way of interpreting the model, when taking differences is necessary, is by viewing it as a regression of involving the differenced variables only and where the error structure follows a stationary and invertible ARMA, That is,

$$\widehat{c}_{i,t} = \boldsymbol{\beta}_i'\widehat{\mathbf{Z}}_{i,t} + e_{i,t}, \tag{7}$$

where $\widehat{c}_{i,t} \equiv \Delta^{d_i}\Delta_s^{D_i}c_{i,t}$ and $\widehat{\mathbf{Z}}_{i,t} \equiv \Delta^{d_i}\Delta_s^{D_i}\mathbf{Z}_{i,t}$ and,

$$\Theta_{i,p_i}(L)\Psi_{i,P_i}(L^s)e_{i,t} = \Gamma_{i,q_i}(L)\Pi_{i,Q_i}(L^s)\varepsilon_{i,t}.$$

Note that, the parameters of interest (i.e. in terms of the example in Eq. (4), the price effects, $\beta_{1,1}$, $\beta_{1,2}$, $\beta_{2,1}$ and $\beta_{2,2}$) are individual coeffcients in the β_i vector and in particular that these are retained under differencing.

3. Confronting the data

The RegARIMA model of sub-period consumption, from the previous section, is now used to analyse price

responsiveness at the intra-day horizon, based on time series consisting of hourly observations of electricity prices and hourly electricity consumption. The price series are the Nord Pool system-or market clearing prices *plus* taxes. Although negative Nord Pool prices sometimes occur, for the present sample this almost never happened (35 out of 43824 hours), and as a short cut we therefore interpolated between the adjacent positive observations. For consumption we consider data corresponding to two consumption levels, the aggregate manufacturing industry consumption of West Denmark and a single anonymous consumer from this industry (henceforth referred to as "Consumer A"). In this way we may also get an idea of the impact of aggregation. Although it is possible that in the aggregate there may be a low degree of "instantaneous price responsiveness", Consumer A could be expected to have at least some possibility of being flexible. This consumption unit, for which the bulk of consumption comes from refrigeration, was picked out since it had the lowest average expenditure per kWh on electricity.

For the aggregate industry we consider the sample January 1st, hour 1 (00-01 AM) 2007 through last hour (11-00 PM) of 2011, a total of 43824 hours (the year 2008 was a leap year). For Consumer A the sample is January 1st, hour 1 (00-01 AM) 2007 through last hour (11-00 PM) of 2010, as some observations were missing for 2011.
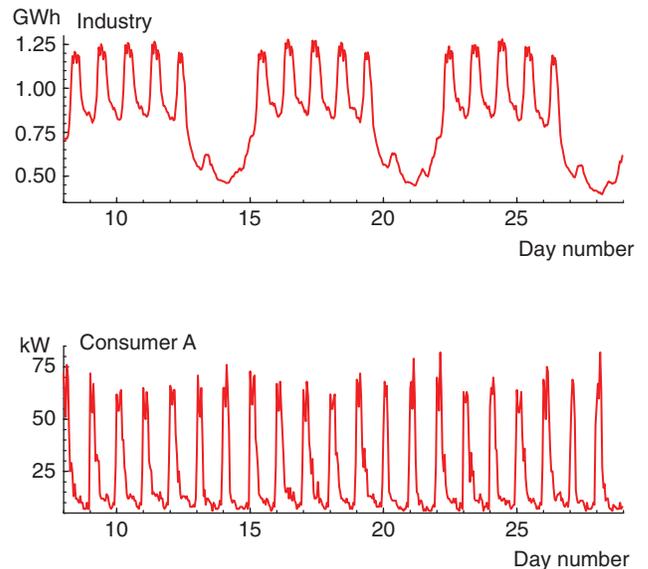

Figure 1: Hourly consumption data, for the aggregate manufacturing industry of West Danmark, in GWh (upper panel) and Consumer A in KW (lower panel) for weeks 2-4 in January 2007.
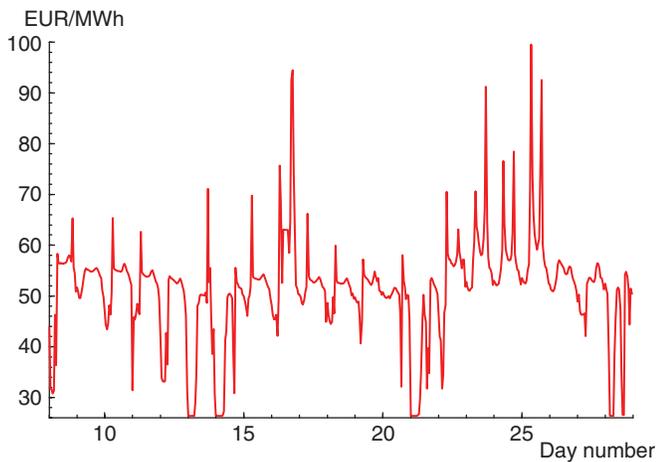
Figure 2: Hourly electricity prices for the aggregate manufacturing industry of West Denmark for weeks 2-4 in January 2007. Units: Euros per MWh (including taxes).

A window (weeks 2-4 of January 2007) of the time series of consumption and prices are shown in Figures 1 and 2. This relatively short window makes the mentioned multi-level seasonality more visible, as compared to a time series plot of the full sample. For example, for the consumption data, and in particular for the aggregate industry, both the weekly and the daily periodicity are strikingly clear.

To have a model that is manageable yet still sufficiently detailed, i.e. time-disaggregated, we have chosen to divide the day into twelve sub-periods each of two hours length. In terms of the regression part of the RegARIMA model, $\boldsymbol{\beta}_i' \mathbf{Z}_{i,t}$, or $\boldsymbol{\beta}_i' \widehat{\mathbf{Z}}_{i,t}$ when differencing is involved (see respectively, Eq. (5) and Eq. (7)), the specification we estimate always includes the 12 price levels from the 12 sub-periods of the current day and the consumption levels for the last 11 sub-periods. In addition, we include a cosine-sine term with frequency, 1/365, to account for the annual seasonality, as resulting from the combined influence from exogenous climatic conditions. Impulse indicators or dummy variables are also included to account for the day (whenever there is no seasonal differencing involved), for the month, for indicating whether the day is a working day, and finally for the industrial holidays.

The estimation is based on the X12-ARIMA module for OxMetrics. See [8] and [17]. For all estimated models we have carried out a residual-based model check.

Following the time series analysis convention this includes a simultaneous assessment of normality and (lack of) serial correlation: The histogram of the residuals was compared against a corresponding normal distribution benchmark. However, since we have many observations the normality assumption is not vital for the statistical results and although the residual distributions are generally rather well-behaved, we accept some non-normality as long as the underlying error distribution can be assumed (approximately) to be symmetrical. In the initial estimations there were often a pronounced degree of skewness (compared to the normal) but in this case this was always due to a (very) limited number of outliers, and hence was taken care of by a few impulse indicators. To check the assumption of no-serial correlation we have considered the significance of the Auto-Correlation-Function (ACF) for the first 65 lags, i.e. a conventional ACF plot with critical values computed under the white noise assumption (see e.g. [16]). As a robustness check all models are also estimated by use of the *automatic* ARIMA modelling implemented in OxMetrics [17].[9] Often, but not always, the automatic choice coincided with our choice of specification. In any case it was clear that the estimated price elasticities and their significance in the tables below were virtually independent of whether we chose the specification manually or automatically.

It is well-known that there is a high degree of correlation between price levels corresponding to the different hours of the day, in particular for adjacent hours. This is an inherent and fundamental problem in light of the regression model and potentially it may lead to finding insignficant price elasticities although the true elasticities are non-zero. However, relative to price levels for adjacent hours, in our application correlations were reduced for two reasons. Firstly, the aggregation into two-hour sub-periods lowered correlation slightly. Secondly, and more importantly the data indicated seasonal unit roots for almost all estimations (see below). Therefore it was necessary to take the 7th difference of the (log-) prices to obtain stationarity and this also lowered correlation to some extent. As a result collinearity is less of a problem than initially expected. It should be mentioned that initially we did experiment with aggregation into a smaller number of longer sub-periods to see whether this would reduce collinearity. It turned out that not much was gained when increasing

---

[9]   See the references in [17].

the length of the sub-periods beyond two hours. For example, dividing the day into four 6-hour sub-periods, the price levels for the adjacent sub-periods were still relatively correlated. To some extent (though less) this was also the case even when dividing the day into two 12-hour sub-periods. Since such aggregation into fewer but longer sub-periods has the cost that potentially interesting information may be hidden by aggregation, we chose to keep the 12 two-hour sub-periods.

Table 1 reports the estimated short-run price elasticities along with their t-ratios, for the aggregate manufacturing industry. The table is read as follows: For example, in the first row corresponding to sub-period 1 consumption, we can see that this responds to the sub-period 1 price, negatively (as expected). The estimated (own-price) elasticity for this sub-period is -0.01 and this is insignificant (t-ratio = 1.16). Note that, own price elasticities on the diagonal are emphasized with grey and t-ratios numerically larger than 2 are marked in bold face. The last column of the table shows the RegARIMA specification for the model for the respective sub-period, and one may note that different sub-periods typically require different

Table 1: Price elasticities for aggregate manufacturing industry consumption in the 12 sub-periods. For each sub-period consumption the first line gives the estimates while the second contains the t-values. Own-price elasticities on the diagonal (emphasized with grey) and t-values numerically larger than 2 are bold faced.

	P1	P2	P3	P4	P5	P6	P7	P8	P9	P10	P11	P12	Model
c1	−0.01	0.00	0.00	−0.01	0.04	−0.01	−0.01	−0.02	0.03	0.03	−0.02	−0.03	RegARIMA
	−1.16	0.19	0.30	−0.83	**3.28**	−0.53	−0.64	−1.72	**2.49**	−1.86	**2.37**	−1.95	$(3,0,0)\times(1,1,1)$
c2	0.00	0.00	0.00	−0.01	0.00	0.01	−0.01	0.01	0.00	0.01	−0.01	0.00	RegARIMA
	−1.30	0.32	0.36	−1.36	−0.26	1.00	−1.38	**2.07**	−0.95	**2.23**	−1.21	0.16	$(0,0,0)\times(0,1,1)$
c3	0.00	−0.01	0.01	0.00	−0.02	0.01	0.01	−0.01	0.00	0.00	0.01	−0.01	RegARIMA
	1.31	−1.59	**2.01**	1.06	**−2.44**	0.82	1.25	−1.85	0.26	−0.32	1.30	−1.21	$(0,0,0)\times(0,1,1)$
c4	0.00	0.00	−0.01	0.01	0.02	−0.02	0.00	0.01	−0.01	0.02	−0.02	0.01	RegARIMA
	0.38	0.67	−0.31	0.84	1.75	−2.15	0.29	1.02	−0.94	1.73	−1.40	1.13	$(2,0,0)\times(0,1,1)$
c5	0.00	0.0	−0.01	0.00	0.01	0.00	−0.01	0.01	0.00	−0.01	0.00	0.00	RegARIMA
	−0.13	1.46	−1.83	−0.43	1.27	0.27	−0.71	0.60	0.78	−1.55	−0.04	0.54	$(3,0,0)\times(0,1,1)$
c6	0.01	0.00	0.00	0.00	−0.01	0.00	0.00	0.01	−0.01	0.01	0.00	0.00	RegARIMA
	2.36	−0.49	−0.52	0.83	−1.40	0.33	−0.19	1.04	−1.20	1.09	−0.25	0.04	$(0,0,0)\times(0,1,1)$
c7	0.01	−0.01	0.01	−0.01	0.01	0.00	−0.01	0.01	0.00	−0.01	0.01	0.00	RegARIMA
	2.11	**−2.60**	**2.06**	−1.30	0.81	−0.45	−1.56	1.69	0.82	−0.93	1.51	−0.59	$(2,0,0)\times(0,1,1)$
c8	−0.01	0.00	0.00	0.01	−0.02	0.01	−0.01	0.01	0.02	−0.01	0.00	0.00	RegARIMA
	−2.04	0.51	0.55	1.96	**−2.63**	1.21	−1.12	1.03	3.67	−1.81	0.04	−0.45	$(4,0,0)\times(1,0,1)$
c9	0.01	0.00	0.00	0.00	−0.01	0.00	0.00	0.00	0.00	0.00	0.00	0.02	RegARIMA
	2.03	−0.83	0.53	0.14	−1.52	−0.34	0.10	0.14	−0.23	−0.46	**2.44**	−1.00	$(0,0,0)\times(1,1,1)$
c10	0.00	0.01	0.00	0.00	−0.01	0.01	−0.01	0.01	−0.01	0.02	−0.01	0.00	RegARIMA
	0.13	1.29	−0.93	0.05	−1.72	1.91	−1.56	0.82	−1.44	**2.81**	−1.05	0.14	$(1,0,0)\times(1,1,1)$
c11	0.00	0.00	−0.01	0.00	0.00	0.00	0.01	0.00	0.00	0.00	0.01	0.00	RegARIMA
	0.07	0.97	−1.70	0.72	−0.11	−0.35	0.54	0.14	−0.05	−0.04	0.65	−0.60	$(1,0,0)\times(1,1,1)$
c12	0.00	0.00	0.00	−0.01	0.02	−0.01	−0.01	0.00	−0.00	−0.01	0.00	0.01	RegARIMA
	−0.68	−0.43	0.73	−1.83	**2.44**	−0.85	−0.75	0.76	0.65	−1.36	0.17	0.92	$(1,0,0)\times(0,1,1)$

specifications although they share common assumptions, such as seasonal differencing, for instance.

Table 2 reports the results for Consumer A and its design is otherwise identical to that of Table 1.

It is clear that both in the aggregate as well as for Consumer A it is hard to find any convincing evidence supporting negative own-price elasticities and positive cross-price elasticities -a hypothesis which could seem reasonable a priori: In general, there are no significant negative own-price elasticities and t-ratios are relatively low throughout the tables. Although insignificant there are five estimates on the diagonal that are negative and in a few cases for Consumer A some positive and moderately

significant cross-price elasticities exist close to the diagonal, which could indicate that some consumption is shifted between sub-periods that are close. An example of this is sub-period 6 for Consumer A, for which there is some, albeit vague, indication that consumption in sub-period 6 goes up if the price in sub-period 7 is high. However, we investigated whether such results would stand out more clearly when removing insignificant regressors, but found that this was indeed not the case.

Although, as mentioned above, taking the 7th difference of the (log-) prices lowered the degree of correlation to some extent, it is inevitable that some of this remained. As a result of this, it could be expected

Table 2: Price elasticities for Consumer A in the 12 sub-periods. For each sub-period consumption the first line gives the estimates while the second contains the t-values. Own-price elasticities on the diagonal (emphasized with grey) and t-values numerically larger than 2 are bold faced.

	P1	P2	P3	P4	P5	P6	P7	P8	P8	P10	P11	P12	Model
c1	0.01	0.02	0.01	−0.04	0.08	−0.10	0.13	−0.05	−0.03	0.05	−0.03	−0.07	RegARIMA
	0.41	0.53	0.15	−1.18	1.51	−1.76	1.86	0.84	−0.77	0.91	−0.39	1.18	(3,0,0)×(1,1,1)
c2	0.00	0.04	0.02	−0.08	0.11	−0.01	−0.12	0.14	0.00	−0.07	0.08	0.00	RegARIMA
	−2.23	0.98	0.40	**−2.39**	**2.16**	−0.19	−1.77	**2.62**	0.00	−1.39	1.30	−0.06	(3,0,0)×(1,1,1)
c3	−0.09	0.00	0.07	−0.06	0.23	−0.19	0.05	0.06	−0.04	0.02	0.03	0.01	RegARIMA
	−1.39	0.03	0.77	−0.96	**2.31**	−1.72	0.36	0.60	−0.51	−0.21	0.21	0.08	(2,0,0)×(0,1,1)
c4	0.03	0.06	−0.08	−0.03	0.09	0.03	−0.02	−0.02	−0.01	0.07	−0.01	0.02	RegARIMA
	0.45	0.96	−1.14	−0.57	1.00	0.33	−0.20	−0.12	0.84	0.78	−0.07	0.25	(3,0,0)×(1,1,1)
c5	−0.06	0.03	−0.03	0.12	−0.11	0.28	−0.28	0.07	−0.03	−0.05	0.09	0.12	RegARIMA
	−1.06	0.41	−0.36	**2.06**	−1.29	**2.99**	**−2.44**	0.79	0.45	0.65	−0.83	1.31	(3,0,0)×(1,1,1)
c6	0.01	0.02	−0.07	0.02	−0.04	−0.08	0.13	−0.02	−0.06	0.08	−0.08	0.05	RegARIMA
	0.48	0.63	−1.76	0.64	−0.92	−1.56	**2.13**	−0.36	−1.74	1.72	−1.31	0.93	(3,0,0)×(1,1,1)
c7	0.00	0.03	−0.04	0.04	−0.01	−0.09	0.11	0.00	−0.03	0.03	0.04	−0.02	RegARIMA
	−0.01	−0.77	−0.94	1.19	−1.32	−1.66	−1.75	−0.08	−0.79	0.68	0.77	−0.39	(2,0,1)×(1,1,1)
c8	0.00	0.01	0.03	0.00	−0.09	0.09	−0.09	0.08	0.00	0.05	−0.03	−0.06	RegARIMA
	−0.06	0.26	0.69	0.01	−1.99	1.73	−1.48	1.68	−0.08	1.00	−0.44	−1.50	(2,0,0)×(1,1,1)
c9	−0.02	0.10	−0.13	0.01	0.01	0.05	−0.14	0.11	0.02	0.00	0.04	−0.03	RegARIMA
	−0.84	**3.11**	**−3.47**	0.38	0.26	1.04	**−2.29**	**2.29**	0.54	−0.10	0.63	−0.58	(2,0,1)×(0,1,1)
c10	0.01	−0.02	0.04	0.00	−0.08	0.05	−0.01	0.03	−0.01	0.04	−0.02	0.04	RegARIMA
	0.32	−0.67	1.26	0.00	**−2.25**	1.28	−1.32	0.95	−0.33	1.10	−0.37	1.02	(1,0,1)×(1,0,1)
c11	0.01	−0.02	0.04	0.00	0.08	0.05	−0.01	0.03	−0.01	0.04	−0.02	−0.04	RegARIMA
	0.32	−0.67	1.26	0.00	**−2.25**	1.28	−0.32	0.95	−0.33	1.10	0.37	−1.02	(1,0,1)×(1,0,1)
c12	−0.03	0.04	0.04	−0.02	0.02	0.04	−0.13	0.10	0.15	−0.22	0.25	−0.14	RegARIMA
	−0.40	0.49	0.37	−0.30	0.15	0.31	−0.86	0.87	1.71	**−0.02**	1.79	−1.15	(2,0,0)×(0,1,1)

that, for example the estimated own-price elasticities on the diagonals could be sensitive to the exclusion of prices from the other sub-periods (i.e. the off-diagonal regressors). For example, given that many of the off-diagonal estimates are marginally insignificant, and hence, supposedly could be excluded, it could be of interest to see whether the estimated own-price elasticities would become negative and significant if all other prices (i.e. from the remaining sub-periods) were removed. We investigated this further but found that this was not the case. Furthermore, for the models with fewer but longer sub-periods, which, as mentioned, we experimented with, the general picture was the same and thus insignificance was a general finding.

Overall it seems safe to claim that, given the adopted modeling approach and the present data, there is no convincing evidence that demand is responding to price changes in the very short run, i.e. at the intra-day horizon.

4. Conclusion and discussion

The idea of Demand Response in power markets has attracted an increasing amount of attention during the last two decades. It has been widely argued that getting consumers to react to short-term variations in electricity prices will improve effciency of electricity markets and assist the integration of renewable production technologies (see e.g. [1], which includes a vast number of relevant references).

A basic premise for this paper has been that demand response is essentially a (short-term) *dynamic* phenomenon, and hence, naturally lends itself to time series modelling. In fact, this dynamicity seems inherent in the usual definition of Demand Response, stating:

"Changes in electric usage by end-use customers from their normal consumption patterns in response to changes in the price of electricity over time, or to incentive payments designed to induce lower electricity use at times of high wholesale market prices or when system reliability is jeopardized."[10]

In this paper we have suggested an econometric model for analysing hourly consumption-and price time series. This allows for multi-level seasonality (i.e. daily, weekly and annual periodicity), which is an inherent characteristic of hourly electricity data [7]. Moreover, the model also takes into account that the information set of agents is updated with new prices only every 24th hour and not every hour, which implies a transformation

of the original series of consecutive hourly observations. Essentially this amounts to treating all hourly observations from a given day as *one* multi-dimensional daily observation. To some extent this approach can be viewed as a generalisation (to a simple demand model involving price and consumption) of that in [6] who analysed price series. Based on this transformation and a general structural model, we derived a dynamic time series model for consumption and prices corresponding to *sub-periods* of the day. In particular, we suggested a linear regression model with multiplicative seasonal ARIMA errors (the RegARIMA), where consumption of a sub-period of the day (e.g. a given hour) is regressed on prices from all sub-periods of that day, and other regressors.

In our empirical analysis of price and consumption data for Danish manufacturing industry for the period 2007-2011, we divided the day into twelve two-hour sub-periods and thus estimated 12 RegARIMA models each of which is a regression of the respective sub-period consumption on prices from all sub-periods of the day and lagged sub-period consumption and prices, in addition to various deterministic variables to take account of annual seasonality, month and day etc..

The overall conclusion to be drawn from the empirical analysis is relatively clear: For the period under study consumer reactions to varying hourly electricity prices have been negligible if not absent. This seems to be the case for both the aggregate manufacturing industry in West Denmark as well as for the single anonymous consumer from this area. This conclusion seems to be rather robust, in the sense that these findings were also obtained based on the automatic ARIMA modeling algorithms as implemented in OxMetrics [17].

Collinearity between prices from adjacent hours could be an issue. On the other hand, in our analysis we made two transformations, that is we aggregate into two-hour periods and, for almost all sub-period models, we take the 7th difference. Both of these transformations reduce collinearity between price regressors. Moreover, we experimented with fewer but longer sub-periods (four 6-hour periods and two 12-hour periods) which a priori could be expected to reduce collinearity further. However, this turned out not to be the case to any notable degree. Moreover, we obtained the same overall conclusion. In future research one could try to pay more attention to this, for example by considering other transformations of the sub-period

---

[10]   This is the definition used by many and it is due to the Federal Energy Regulatory Commission.

variables, such as ratios of sub-period consumption and prices instead.

The RegARIMA is rather general, in that it nests a number of applied time series models in the literature, i.e. seasonal and non-seasonal AR, MA, regressions with white noise errors etc. [7, Section 3.4]. However, other classes of time series models could also be applied with advantage in future research. Examples include models with thresholds for (symmetric or asymmetric) adjustment of consumption to price changes, models with other non-standard (e.g. heavy-tailed or ARCH-type) error distributions. Furthermore, as extreme observations often occur in electricity time series data, a thorough analysis of the influence from such observations (or small groups of) on estimation, could also uncover potentially interesting results. We did not attempt to carry out any of such econometric analyses since that would be beyond the scope of this paper and, in our view, deserves a thorough treatment elsewhere.

There are a number of possible explanations for the apparent lack of short-term response. First, it is possible that many of the industrial consumers are in fact too small and/or have too low energy intensity in their production to benefit from being more flexible. Indeed, looking at Danish industrial consumers many companies are small/medium sized companies with a relatively low energy intensity. That is, the total electricity bill is relatively low and constitutes only a minor share of total production costs. In addition, as taxes and grid payments are fixed per kWh and comprise about one half of the electricity bill, the gain from demand response may be quite limited for many companies. Secondly, the hourly variations in electricity prices for these data may simply be too small for the potential gains to be of a significant order of magnitude. Third, information costs implied by monitoring hourly prices may be perceived as relatively large. Finally and perhaps most importantly, for many industrial consumers the costs associated with adjusting production are likely to be relatively large compared to what may be saved on the electricity bill when moving consumption to hours of cheap electricity. In particular, for many companies changing electricity consumption with short notice is likely to imply idle production capacity and workforce, lost production or reduced product quality.

In addition to these explanations we also emphasize the fact that our findings do not exclude the possibility of demand response, as such. One has to add a little nuance here. What we investigate is an hour-by-hour, i.e. "continuous" response. However, it may well be the

case that some consumers, although paying for hourly consumption, still choose not to follow prices in a continuous manner due to the costs associated with doing so. For example, agents may instead adhere to a pre-specified rule, such as to consume less during the day and more during the night when the average price is much lower. Supposedly such a rule is only updated in the very long run or when large prices changes take place. This could explain our findings concerning the single consumer (Consumer A). Indeed, for this consumer, who was evaluated to have both the incentive and the ability to be flexible, simply plotting the time series of prices and consumption together against time, clearly suggests that this may well be the case. This is done in Figure 3.

In the future with increased production from renewable production technologies the volatility of prices and the need for flexibility is expected to increase. Incentives for demand response may therefore increase, but other changes may be required for demand response to play a significant role. For example, grid payments and possible taxes may be changed to follow hourly market prices on electricity. This increases incentives for flexibility but may conflict with market efficiency. Furthermore, technology and automatic control of consumption will decrease information and monitoring costs which currently could seem to hamper demand response. However, for mass market (e.g. household appliances) control technology is required to be cheap and acceptance of automated control may be a problem. For larger consumers automated control of part of the consumption (e.g. heating, cooling and pumping) may be acceptable, but potential gains are limited by production schedules and product quality.
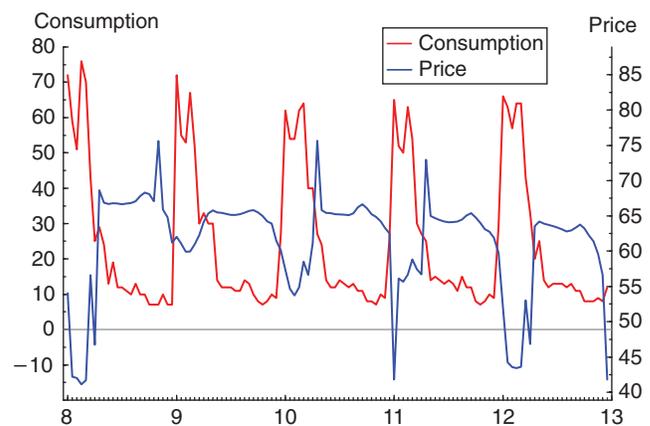

Figure 3: Hourly consumption (first axis) and prices (second axis) for Consumer A for 5 days in january 2007.

To increase demand flexibility, focus areas should be consumers with a large potentially flexible consumption and the development of cheap information and control technologies. Many argue that household appliances like freezers, coolers and water heaters may supply cheap flexibilities as the technologies may be shut off for shorter periods without notable consequences for the consumer. However, today household consumers face fixed short-term prices. Still, exposing households to hourly prices, to harvest this flexibility, realistically automatic control is required. In addition, as the savings in the electricity bill are minor, control technologies have to become very cheap, and acceptance of automated control may be a barrier.

Lastly, we would like to emphasize that, at least to some extent, the present empirical analysis has also served the purpose of illustrating an application of our econometric model. It is also clear that more data sets have to be scrutinized in order to provide a more solid basis for making any generalisations. Nevertheless, we believe that our econometric approach is relatively general, in that it allows for a large number of interesting combinations of number and length of sub-periods. Despite the fact that the price data we have analysed supposedly do not exhibit the sufficient amount of variation to induce behavioural responses in the short run, it is therefore our hope that others will apply our method to data sets with more variation, the latter of which most likely will be widely available in the future.

Acknowledgement

References

[1] US Department of Energy, Benefits of Demand Response in Electricity Markets and Recommendations for Achieving Them, Tech. rep., US Department of Energy (2006).

[2] D. W. Caves, L. R. Christensen, J. A. Herriges, Consistency of residential customer response in time-of-use electricity pricing experiments, Journal of Econometrics, Elsevier 26 (1–2) (1984) 179–203. URL *http://EconPapers.repec.org /RePEc:eee:econom :v:26:y:1984:i:1-2:p:179–203*

[3] A. Faruqui, S. S. George, The value of dynamic pricing in mass markets, The Electricity Journal 15 (6) (2002) 45 -55.

[4] C. S. King, S. Chatterjee, Predicting California Demand Response, Public Utilities Fortnightly (2003) 27-32.

[5] A. Faruqui, S. George, Quantifying customer response to dynamic pricing, The Electricity Journal 18 (4) (2005) 53–63. URL *http:// EconPapers.repec.org /RePEc:eee:jelect:v: 18:y :2005 :i:4:p:53–63*

[6] F. A. Wolak, Market Design and Price Behavior in Restructured Electricity Markets: An International Comparison, in: Deregulation and Interdependence in the Asia-Pacific Region, NBEREASE Volume 8, NBER Chapters, National Bureau of Economic Research, Inc, 2000, pp. 79-137. URL *http://EconPapers.repec.org/RePEc:nbr: nberch:8478*

[7] R. Weron, Modeling and Forecasting Electricity Loads and Prices: A Statistical Approach, Wiley Finance, 2006. URL *http://EconPapers.repec.org/RePEc:wuu:hsbook:hsbook0601*

[8] D. F. Findley, B. C. Monsell, W. R. Bell, M. C. Otto, B. C. Chen, New capabilities and methods of the x-12-arima seasonal-adjustment program, Journal of Business and Economic Statistics 16 (2) (1998) 127-52. URL *http:// EconPapers. repec.org/RePEc:bes:jnlbes:v:16:y:1998:i:2 :p:127–52*

[9] R. Huisman, C. Huurman, R. Mahieu, Hourly electricity prices in day-ahead markets, Energy Economics 29 (2) (2007) 240–248. URL *http://EconPapers.repec.org/ RePEc:ems :eureri:8289*

[10] R. Ramanathan, R. Engle, C. W. J. Granger, F. Vahid-Araghi, C. Brace, Short-run forecasts of electricity loads and peaks, International Journal of Forecasting 13 (2) (1997) 161–174.

[11] C. Fezzi, D. Bunn, Structural analysis of electricity demand and supply interactions, Oxford Bulletin of Economics and Statistics 72 (6) (2010) 827–856.

[12] M. D. Intriligator, Economic and econometric models, in: Z. GrilichesâĂ, M. D. Intriligator (Eds.), Handbook of Econometrics, Vol. 1, Elsevier, 1983, Ch. 3, pp. 181–221. URL *http://EconPapers.repec.org/RePEc:eee:econhb:1*

[13] R. Davidson, J. MacKinnon, Econometric theory and methods, Oxford Univ. Press, 2004.

[14] H. Hamilton, Time Series Analysis, Princeton University Press, 1994.

[15] S. Johansen, Likelihood-Based Inference in Cointegrated Vector Autoregressive Models, Advanced Texts in Econometrics, Oxford University Press, Oxford, 1996.

[16] G. Box, G. Jenkins, G. Reinsel, Time Series Analysis: Forecasting and Control, Wiley Series in Probability and Statistics, Wiley, 2008.

[17] J. Doornik, D. Hendry, PcGive 14, OxMetrics 7, Timberlake Consultants, 2013.

[18] D. F. Hendry, Dynamic Econometrics, Oxford University Press, 1995.