Uncovering the nonlinear predictive causality between natural gas and electricity prices[1]

Jorge M. Uribe [a,b]

Montserrat Guillen[a,c]

Stephania Mosquera-López [d]

[a] *Riskcenter, University of Barcelona, Spain*

[b] *Department of Economics, Universidad del Valle, Colombia*

[c] *Department of Econometrics, University of Barcelona, Spain*

[d] *School of Industrial Engineering, Universidad del Valle, Colombia*

This version February 8th 2018

Abstract

We measure the directional predictability between electricity and natural gas prices at different quantiles of the two price distributions. This enables us to uncover significant nonlinearities in the relationship that feature the markets of gas and electricity in New England and the Pennsylvania-New Jersey-Maryland interconnected market. We show that there is a double causality from gas to electricity and the other way around, which is larger when the observed prices in the market are higher. In general these effects are considerably high for the median-up ranges of the two prices. That is, when both electricity and gas prices are recorded above the median market prices. The feedback effect from electricity to gas is stronger for New England where 50% of power generation mix consists of natural gas-fired plants, than for Pennsylvania-New Jersey-Maryland, where 24% of the generation mix relies on natural gas sources.

Keywords: Natural gas; Electricity; Directional predictability; Quantiles; Cross-quantilograms.

[1] The support received from the Spanish Ministry of Science FEDER ECO2016 -76203-C2-2-P is acknowledged. The second author thanks ICREA Academia.

1. Introduction

We carry out a systematic examination of the relationship between natural gas and electricity prices in the Pennsylvania-New Jersey-Maryland (PJM) and the New England (NE) interconnection markets, by the means of recently proposed cross-quantilogram functions (Han et al., 2016). Our approach enables us to test for directional predictability running from natural gas to electricity prices and from electricity to gas, not only at different lags, but also at different conditional quantiles of the distribution of the two variables. In order to understand the dynamics of energy markets featured by strong seasonalities both in demand and supply, we propose to test for directional predictability between these variables while we consider nonlinearities across the quantiles of the energy price distribution. Hence, our results are relevant for the optimal design of hedging mechanisms by generators and consumers of electricity and natural gas, and for the construction of an optimal policy framework under which the market may operate smoothly by mitigating, for instance, episodes such as the California energy crisis extended from May 2000 to June 2001.

Our analysis builds a bridge between two different branches of the recent literature in energy economics. The first one has studied and measured directional predictability between electricity and natural gas prices (see for instance Woo et al., 2006; Brown & Yücel, 2008; Chae et al., 2012; Nakajima & Hamori, 2013; Alexopoulos, 2017) and the second one has analyzed the price dynamics of energy prices across different fragments of their distribution (see Bunn et al., 2016; Hagfors et al., 2016; Mosquera-López et al., 2017). By so doing we are able to characterize significant nonlinearities that describe the relationship between natural gas and electricity prices. To the best of our knowledge, we are the first in describing these nonlinear features, and therefore the first that can give some advice to market practitioners and regulators that explicitly considers the price level recorded in the market on a given date. As mentioned above, it turns out that in terms of predictability and contemporaneous association, the joint gas and electricity market behaves remarkably different when prices are relatively high or relatively low. For example, we show that there is higher predictability from natural gas to electricity when the two series of prices are at a high quantile (i.e. their 90th percentile), and that this predictability reduces considerably for other quantiles, particularly for the lowest quantiles of the energy prices (i.e. its 10th percentile).

We also present significant evidence of a double causality between the two prices, which has been neglected in previous studies that focus only on the mean-to-mean effect. i.e. how a change in the mean of the first variable predicts a change in the mean of the second variable (see Brown & Yücel, 2008 and Alexopoulos, 2017), or even in studies that document a mean-to-variance effect (Nakajima & Hamori, 2013). This double causality can be understood following Woo et al., (2006) as a consequence of a *supply-push* effect and a *demand-pull* effect. The former refers to the well-documented impact of gas on electricity. Natural gas prices affect market clearing and price formation in electricity markets in two ways as stressed out by Alexopoulos (2017). The first one emerges because gas based power producers are frequently the last ones to be included in the merit-order curve and therefore they determine the wholesale electricity rates, which also affects retail electricity rates provided by Load Serving Entities (LSE). In the PJM and NE markets studied here, indeed, the generation mix of electricity includes 24% and 50% of power that is produced using natural gas power sources, respectively[2]. Second, LSE may own directly gas power stations and, in this case, they implement a direct pass through with automatic mechanisms of the unexpected fuel costs to electricity consumers. The other direction of the causal relationship, namely from electricity to gas, is less understood and has been frequently overlooked in the literature. Woo et al., (2006)

[2] Generation mix data source: https://www.eia.gov/electricity/data.php#generation

refers to it as a demand-pull effect that operates as follows. Suppose that a random increase in electricity demand triggered by cold weather is observed. Then, *ceteris paribus,* such an increment in the demand will translate into a wider spark-spread between electricity prices minus the fuel cost of a natural gas-fired power plant. The larger spread will raise the demand for natural gas not only by increasing the willingness to pay for it, but also by inducing less efficient plants to start operating. As has been explained by Woo et al., (2006) this in turn will manifest in higher bids for spot gas in bilateral trading and higher realized natural gas prices. These authors also highlight that following this general mechanism, depending on the elasticity of natural gas supply, an increase in the demand for electricity may translate one-to-one into higher natural gas prices, which in turn may generate a feedback effect on electricity prices, leading to a possible scaling up of the original shocks that may endanger the operation of the two markets. Here it becomes fundamental some previous insights from the recent literature that point out prominent nonlinearities in the price distribution of electricity, which condition on fundamental factors such as weather and market variables (Bunn et al., 2016; Hagfors et al., 2016; Mosquera-López et al., 2017). For instance, Mosquera-López et al., (2017) show that weather factors induce a significant and nonlinear influence on electricity prices, which varies depending on the conditional percentile of the prices. They also document that, as expected, the main influence of weather occurs at the tails of the electricity price distribution, where abnormally high and low prices are recorded. We expect these sorts of nonlinearities to manifest themselves in an asymmetric reaction of electricity prices following changes in natural gas prices and vice versa. Thus, we study by the first time this possibility here, and we find that indeed this is the case. The predictive-causality between electricity prices and natural gas prices depend at a large extent on the observed market prices, and therefore, vary greatly according to the quantile prices of both electricity and gas. We also find that the intricate double causality that characterizes the operation of the two markets seems to be related to the extent to which electricity generators rely on natural gas power sources. Thus, the higher the generation mix of electricity is composed by natural gas, the higher the predictive causality from natural gas to electricity at all the quantiles of the price distribution. We support this latter assertion based on a comparison of our main results for PJM and NE.

We show that the predictive causality goes in the two directions, from gas to electricity and from electricity to gas. Nevertheless, the relationship changes remarkably conditioning on different quantiles of the price distribution. In both markets the strongest directional predictability is recorded when the electricity and gas prices have reached their 90^{th} percentile, which is precisely when market impairments are more likely. In PJM the cross-correlation between the quantile-hit functions lie between 20% and 30%, from gas to energy, and it is around 15% from electricity to gas, at the 90^{th} percentile of the series (although it only lasts for one week in the latter case, in contrast with the former that remains significant as long as two months). In the New England market, which in contrast to PJM is described by a mix of primary sources that heavily relies on natural gas, these numbers range in both directions, between 40% and 60%, and the relationship lasts approximately 40 days.

The rest of this paper is organized as follows. In section 2 we present our main methods. In section 3 we describe our data and data sources. Our results are reported in section 4, regarding the PJM and the NE markets. In section 5 we conclude, summarize our main findings and discuss policy implications for energy markets, their regulators and practitioners.

2. Methodology

Our methodology relies on the work by Han et al. (2016), who proposed the cross-quantilogram as a measure of serial dependence between two series at different conditional quantile levels, and an application of the Box-Ljung Q-statistic to test for directional predictability between the two series. Thus, we followed closely their presentation and notation in this section. Notice that the cross-quantilogram is a generalization of the quantilogram proposed by Linton and Whang (2007), which is indeed equivalent to the former in the special case of a single series.

2.1. *The cross-quantilogram*

Let $\{(\boldsymbol{y}_t, \boldsymbol{x}_t): t \in \mathbb{Z}\}$ be a stationary price time series with $\boldsymbol{y}_t = (y_{1t}, y_{2t})' \in \mathbb{R}$ and $\boldsymbol{x}_t = (x_{1t}, x_{2t}) \in \mathbb{R}^{d_1} \times \mathbb{R}^{d_2}$, where $x_{it} = \left[x_{it}^{(1)}, \ldots, x_{it}^{(d_i)}\right]' \in \mathbb{R}^{d_i}$ and $d_i \in \mathbb{N}$ for $i = 1{,}2.$ We denote $F_{y_i|x_i}(\cdot|x_{it})$ to the function of the series y_{it} given x_{it}, and $f_{y_i|x_i}(\cdot|x_{it})$ to the corresponding density. In this context, $q_{it}(\tau_i) = \inf\{v: F_{y_i|x_i}(v|x_{it}) \geq \tau_i\}$ for $\tau_i \in (0{,}1)$, and for $i = 1{,}2.$ The cross quantilogram is a measure of serial dependence between two events $\{y_{1t} \leq q_{1,t}(\tau_1)\}$ and $\{y_{2t-k} \leq q_{2t-k}(\tau_2)\}$ for any pair of $\tau = (\tau_1, \tau_2)'$ and for an integer k. It can be understood as the cross-correlation of the quantile-hit processes given by $\{1[y_{it} \leq q_{it}(\cdot)]\}$, and is defined by the following equation:

$$\rho_\tau(k) = \frac{E[\psi_{\tau_1}(y_{1t}-q_{1t}(\tau_1))\psi_{\tau_2}(y_{2t-k}-q_{2t-k}(\tau_2))]}{\sqrt{E[\psi_{\tau_1}^2(y_{1t}-q_{1t}(\tau_1))]}\sqrt{E[\psi_{\tau_2}^2(y_{2t-k}-q_{2t-k}(\tau_2))]}}, \quad (1)$$

for $k = 0, \pm 1, \pm 2, \ldots$, where $\psi_a(u) \equiv 1[u < 0] - a$. It is a well-defined statistic even for processes with infinite moments and it is invariant to any strictly monotonic transformation applied to both series, including taking logs.

The sample counterpart of equation 1, which relies on a linear quantile regression as proposed by Koenker and Bassett (1978) to estimate the conditional quantile functions, $\hat{q}_{it}(\tau_i)$, is given by:

$$\hat{\rho}_\tau(k) = \frac{\sum_{t=k+1}^{T} \psi_{\tau_1}(y_{1t} - \hat{q}_{1t}(\tau_1)) \psi_{\tau_2}(y_{2t-k} - \hat{q}_{2t-k}(\tau_2))}{\sqrt{\sum_{t=k+1}^{T} \psi_{\tau_1}^2(y_{1t} - q_{1t}(\tau_1))} \sqrt{\sum_{t=k+1}^{T} \psi_{\tau_2}^2(y_{2t-k} - q_{2t-k}(\tau_2))}}. \quad (2)$$

Notice that the cross-quantilogram considers dependence in terms of the direction of deviation from conditional quantiles and therefore measures the directional predictability from one series to another. For this reason what is series 1 and what is series 2 matters in our empirical implementation in which we considered the two cases for natural gas and electricity prices. $\hat{\rho}_\tau(k)$ lies between -1 and 1, and naturally $\hat{\rho}_\tau(k) = 0$ corresponds to the case of non-directional predictability.

2.2. *Q-test for directional predictability*

In the case in which one is interested in testing the null hypothesis $H_0: \rho_\tau(1) = \cdots = \rho_\tau(p) = 0$ against the alternative that $\rho_\tau(k) \neq 0$ for some $k \in \{1, \dots, p\}$, that is, in testing the directional predictability of events up to p , Han et al. (2016) propose to use a traditional Box-Pierce statistic for small samples (Box-Ljung) given in this context by:

$$\hat{Q}_\tau^p \equiv T(T+2) \sum_{k=1}^{p} \hat{\rho}_\tau^2(k)/[T-k], \quad (3)$$

2.3. *The stationary bootstrap*

Han et al., (2016) propose as well the stationary bootstrap of Politis and Romano (1994) to approximate the null distribution of the cross-quantilogram and the Q-statistic presented above, while avoiding the dependence on nuisance parameters of the asymptotic distribution. The stationary bootstrap is a method for block bootstrapping with blocks of random lengths.

Let $\{K_i\}_{i\in\mathbb{N}}$ be a sequence of iid random variables, which are drawn from a discrete uniform distribution on $\{k+1,\ldots,T\}$, and are independent on the original data and a sequence $\{L_i\}_{i\in\mathbb{N}}$ of random block lengths. B_{K_i,L_i} represents the blocks of length L_i that are used to generate the bootstrapped samples $\{(\boldsymbol{y}_{tk}^*, \boldsymbol{x}_{tk}^*)\}_{t=k+1}^{T}$. These (re)samples are then use to estimate the conditional quantile function $q_{it}^*(\tau_i)$ for each $i = 1,2$. The cross-quantilogram based on the stationary bootstrap resample is given by:

$$\hat{\rho}_\tau^*(k) = \frac{\sum_{t=k+1}^{T} \psi_{\tau_1}(y_{1t}^* - \hat{q}_{1t}^*(\tau_1))\psi_{\tau_2}\left(y_{2t-k}^* - \hat{q}_{2t-k}^*(\tau_2)\right)}{\sqrt{\sum_{t=k+1}^{T} \psi_{\tau_1}^2\left(y_{1t}^* - \hat{q}_{1t}^*(\tau_1)\right)}\sqrt{\sum_{t=k+1}^{T} \psi_{\tau_2}^2\left(y_{2t-k}^* - \hat{q}_{2t-k}^*(\tau_2)\right)}}, \qquad (4)$$

then the bootstrap is used to construct a confidence interval for each statistic of p cross-quantilograms $\{\hat{\rho}_\tau(1),\ldots,\hat{\rho}_\tau(p)\}$ for a finite positive p (lags). The technique is also followed to construct a confidence interval for the Q-statistic presented in equation 3. See Han et al., (2016) for a formal proof of the validity of the procedure and for further technical clarifications.

The authors also advanced a self-normalized version of the cross-quantilogram and the Q-statistic, following the works by Lobato, (2001) and Shao (2010). We also considered this alternative in our empirical implementation but our results remain unchanged compared to the case when the inference was carried out using the stationary bootstrap, which has been just described, thus in the sake of space we do not report them. Moreover Han et al., 2016 document a better performance in term of power of the stationary bootstrap over the competing alternatives for small samples in their Montecarlo simulations.

3. Data

The U.S. wholesale electricity market has traditionally regulated areas (vertically-integrated utilities responsible for the generation, transmission and distribution systems) and restructured competitive markets. The latter ones, named as Regional Transmission Organizations (RTO)/Independent System Operators (ISO), are responsible for serving two-thirds of the country's load, and wholesale market participants can bid or offer electricity generation (see the colored areas in Figure 1). The largest ISO operating in the U.S., and one of the largest in the world, is the Pennsylvania-New Jersey-Maryland (PJM) interconnection with an installed capacity of 165 GW, a geographical footprint of 243,417 squared-miles, and serves 13 states (Delaware, Illinois, Indiana, Kentucky, Maryland, Michigan, New Jersey, North Carolina, Ohio, Pennsylvania, Tennessee, Virginia, West Virginia and the District of Columbia) with a population of 65 million people[3]. PJM interconnection was founded in 1927 with a pool of only three utilities in Pennsylvania and New Jersey. In 1956, Maryland entered the interconnected market, and by 1996 PJM became a fully functioning ISO.

Given the importance of PJM in the U.S. wholesale electricity market, we chose this market for our analysis of the directional predictability between electricity and natural gas prices. Hence, natural gas in PJM's generation mix stands only for the 24%, we also carry out or study for the ISO functioning in New England (ISO-NE), where natural gas is the principal source and represents the 50% of the system's supply. We use the daily spot electricity and natural gas prices of the hubs in PJM and New England markets published by the Intercontinental Exchange (ICE), and republished by the U.S Energy Information Administration (EIA). The frequency of the prices is daily, and the sample period starts on March 2014, and ends on December 2017.

[3] Data source: https://www.ferc.gov/market-oversight/mkt-electric/pjm.asp

Figure 1. Price hub locations for wholesale electricity and natural gas reported by Intercontinental Exchange

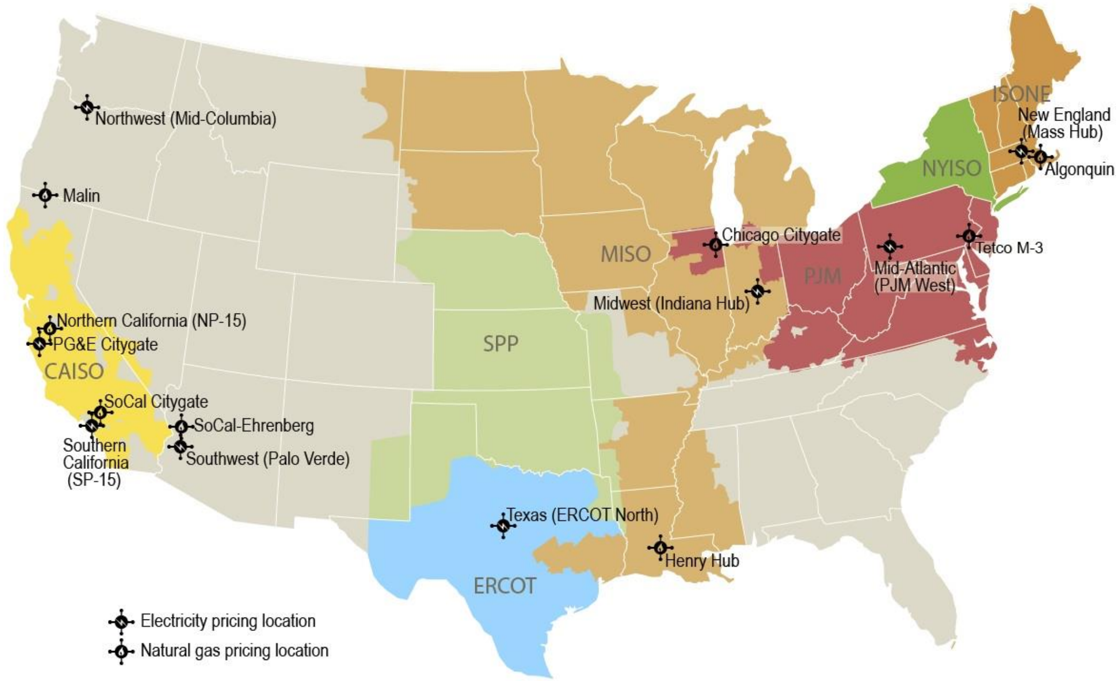

Note: the colored areas correspond to Regional Transmission Organizations (RTO)/Independent System Operators (ISO). Retrieved from U.S Energy Information Administration (EIA).

Table 1 presents the descriptive statistics for the prices of the two regions, and Figure 1 depicts them. PJM WH Real Time Peak is the ICE electricity product name for the electricity hub PJM West, and TETCO-M3 is the name of the product from the natural gas hub TETCO-M3, both from the region of PJM. Nepool MH DA LMP Peak is the ICE electricity product name for the electricity hub Mass Hub, and Algonquin Citygates is the product name from the natural gas hub Algonquin of the New England region. On average Nepool electricity prices are higher than PJM's prices, which accounts for PJM being a bigger more competitive market. In the case of natural gas, the same dynamics is found. Moreover, New England's electricity and natural gas prices tend to be more volatile than PJM's prices, but they present similar dynamics in terms of peaks in the market.

Table 1. Descriptive statistics wholesale electricity and natural gas spot prices for PJM and New England markets

Region	*PJM*		*New England*	
	PJM	**TETCO-M3**	**Nepool**	**Algonquin Citygates**
Mean	39.51	2.43	45.61	4.49
Std. Dev.	15.77	1.98	26.14	4.11
Minimum	22.70	0.38	16.00	0.78
10th percentile	28.00	1.21	25.62	1.82
50th percentile	36.04	1.92	38.77	3.26
90th percentile	52.86	3.76	67.79	8.13
Maximum	231.54	20.65	199.73	52.22
Observations	946	946	656	656

Figure 2. Wholesale electricity and natural gas spot prices for PJM and New England markets

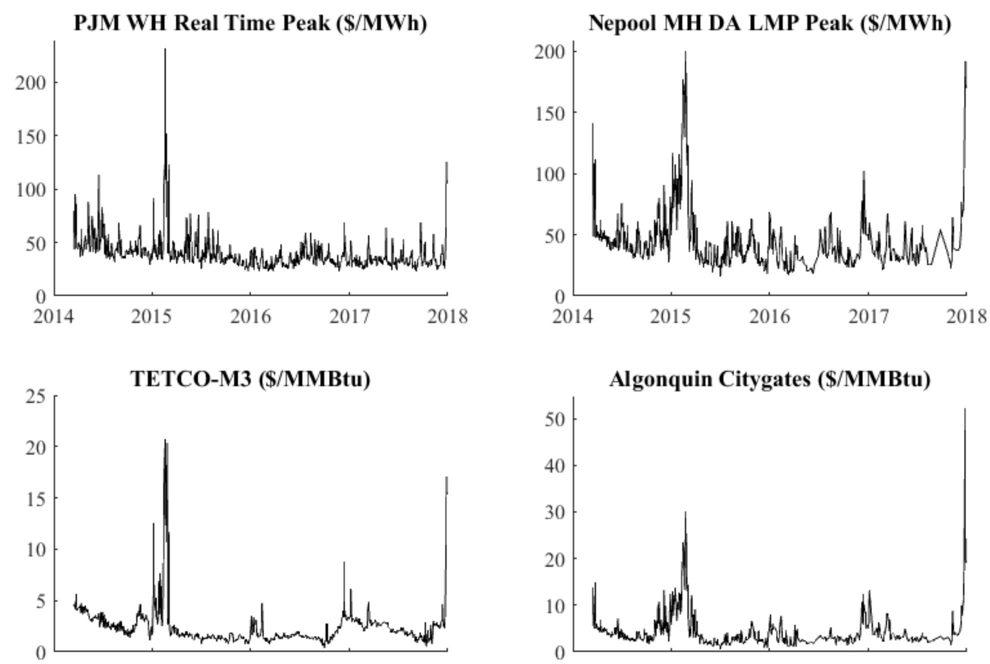

Note: own elaboration with data from the Intercontinental Exchange (ICE) republished by the U.S Energy Information Administration (EIA).

4. Results

1.1. *Results for the Pennsylvania-New Jersey-Maryland interconnection*

Figure 3 shows the estimated cross-quantilogram that measures directional predictability at different quantiles and different lags between gas and electricity prices at the PJM market, with the causality (in the Granger sense) running from the series of natural gas to the series of electricity. The results vary depending on the analyzed conditional quantile, both in the predicted variable (i.e. electricity power), which is presented in the rows of the figure, and the predicting variable (i.e. natural gas), which is shown in columns. Figure 4 depicts the Q-statistics that test the null of non-predictability from gas to electricity at different cumulated lags and that condition on different quantiles also at the PJM market. The larger pockets of predictability appear when electricity prices are near to their median range (50^{th} percentile), or when both electricity and gas prices are at their highest quantiles (90^{th} percentile in our estimations). In both cases the cross-correlation among the quantile hit functions ranges approximately between 15% and 25%, and it can even increase above 30% when both electricity and gas prices reach the 90^{th} percentile. It is also noticeable that predictability is considerably lower when we focus on the lowest percentiles of the electricity prices (10^{th} percentile), being bellow 15% by general rule and statistically insignificant in most of the cases. Additionally, we also observe that the predictability from gas to electricity seems to increase according to the quantile of the natural gas prices. That is, it turns out that the higher the prices of natural gas (for example due to abnormally low temperatures or a relative scarcity of the resource) the more these prices can be used to predict future electricity prices. This is true disregarding the electricity quantile, but it is especially true for prices of electricity recorded above its own median.

Figure 3. Cross-quantilogram to measure the directional predictability from natural gas to electricity prices: PJM

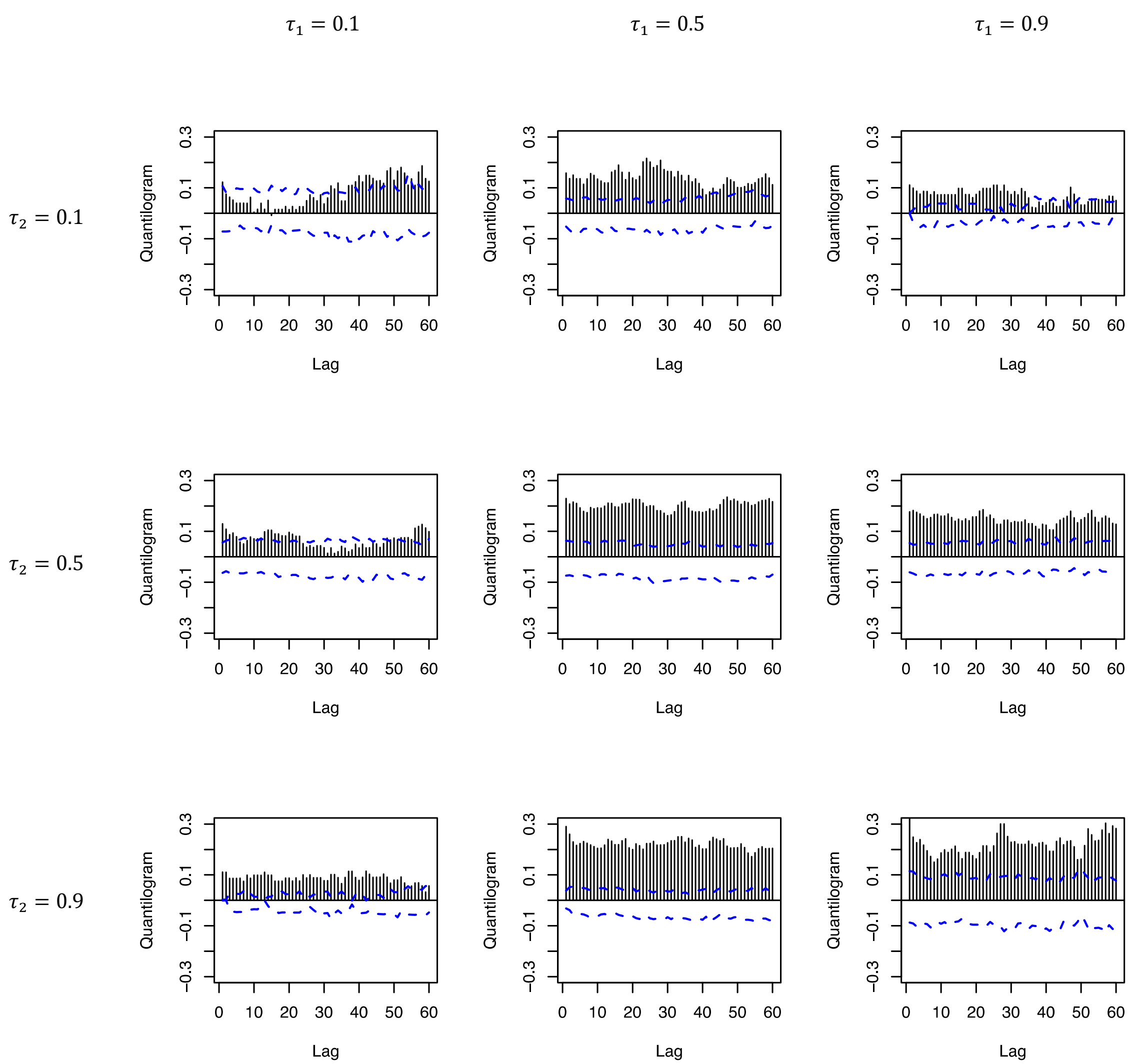

In the figure are presented the sample cross-quantilograms $\hat{\rho}_\tau(k)$ for $k = 1 \dots 60$ and $\tau_i = (0.1, 0.5, 0.9)$ for $i = 1,2$. The cross-quantilogram measures the directional predictability from natural gas to electricity prices, in PJM. In the rows are sorted the quantiles of the natural gas prices, while in the columns of the electricity prices. The bars represent $\hat{\rho}_\tau(k)$, while the dotted lines are confidence intervals at 95% of confidence constructed by the stationary bootstrap centered at zero with $B = 1000$ number of repetitions.

In terms of the lag structure, which measures how the quantile correlation changes depending on the number of lags separating electricity and natural gas prices, we also document significant asymmetries conditioning on different quantiles. For example, when both electricity and natural gas prices are low (namely at their 10^{th} percentile), there is not predictive power running from gas to electricity approximately up to the lag 30. That is, the cross-qauntilogram becomes significant in this case only after days 30-35. On the contrary, when the price of electricity is relatively high (90^{th} percentile) and the gas is still cheap, the pattern is just the opposite: some predictability arises from natural gas to electricity up to lag 30 and it disappears afterwards. For the ranges in which both electricity and natural gas prices are above their median the cross-qauntilogram displays a clear persistence across all lags.

Turning our attention to Figure 4, we confirm the causality (in terms of predictability) running from natural gas prices to electricity prices, especially for the quantiles of electricity above its own median. In these cases the Q-statistic is always above the bootstrapped confidence interval at the 95% level of confidence. When the price of electricity is relatively low (10^{th} percentile) the predicting power of natural gas prices is considerably reduced, particularly when its own prices are bellow its median, in which the Q-statistic signals non-directional predictability at most of the lags. Once again it is possible to note that the highest Q-statistics, informing about a large predictability from gas to electricity, arise when both series are recorded at their respective 90^{th} percentiles. Remarkably, these high prices naturally correspond to situations of relatively scarcity of both resources, perhaps due to a reduction in the temperature and other weather effects that exacerbate the energy demand, or produce a shortage of the energy supply, and push up the market prices.

Figure 4. Box-Ljung test and directional predictability from natural gas to electricity prices: PJM

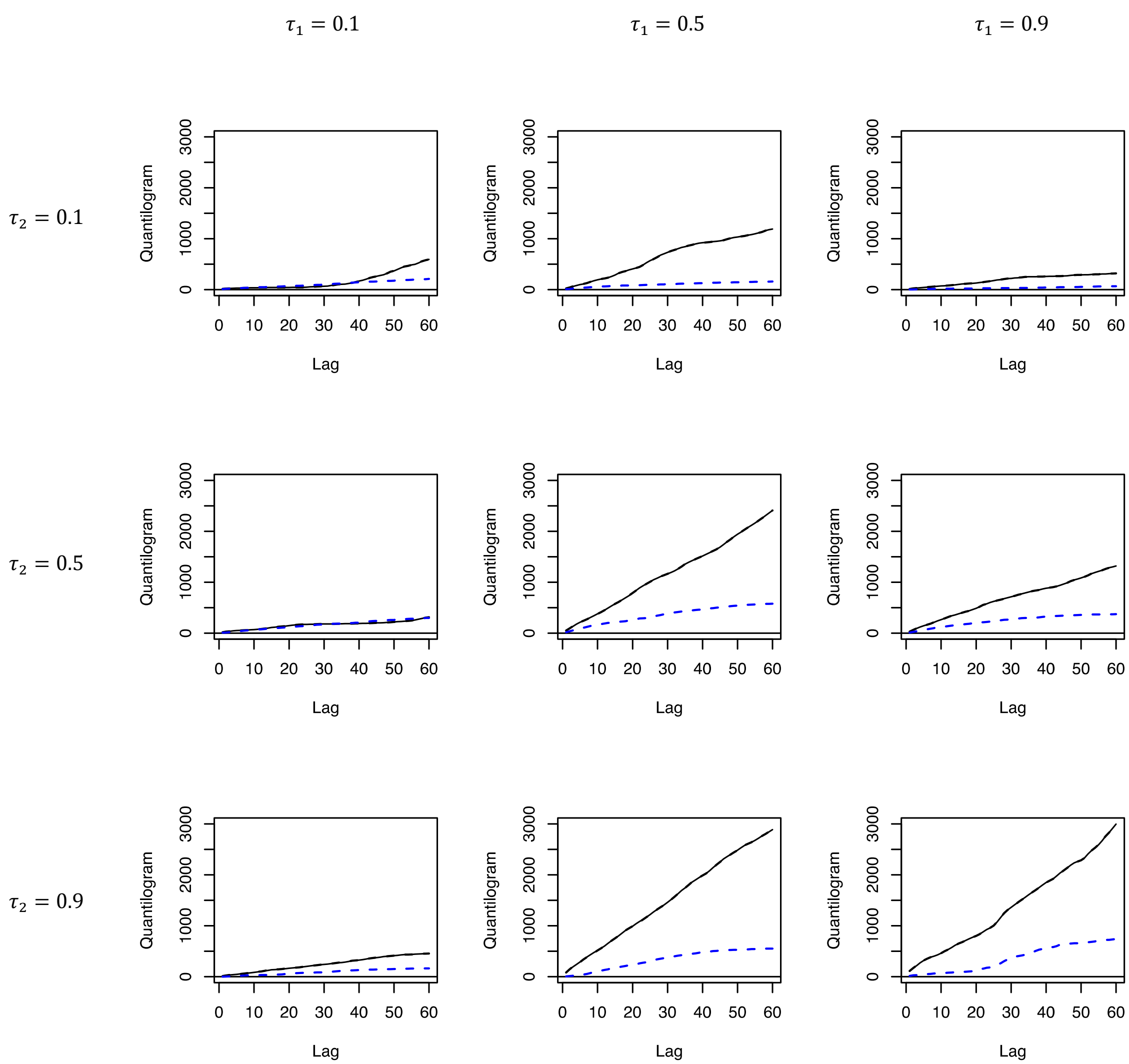

In the figure are presented the sample $\hat{Q}_{\tau}^{(p)}$ for $p = 1 \dots 60$ and $\tau_i = (0.1, 0.5, 0.6)$ for $i = 1,2$. The solid lines are the Q-Box-Ljung statistics, while the dotted lines are confidence intervals at 95% of confidence, constructed by the stationary bootstrap centered at zero with $B = 1000$ number of repetitions.

Now we focus on the predictability from electricity to natural gas prices. As explained in the introduction, this is an often-overlooked relation that we find, however, to be crucial. In Figure 5, we present the cross-quantilogram for the PJM market, with the causality running from electricity to natural gas prices, and in Figure 6 we present the respective Q-statistics. When we focus on the quantiles close to the median of either electricity or natural gas prices, we document a significant causality (in the Granger sense) from electricity to natural gas. This predicting power of electricity on natural gas reduces considerably for the extreme quantiles of both energy prices. That is, when we consider the left-top scenario depicted by Figure 3, in which both electricity and natural gas prices are relatively low (10^{th} percentile), or the scenarios on the top-right and bottom-left, in which either electricity or natural gas a relatively low-priced, the predictive power of electricity on natural gas almost disappears (or it is rendered as statistically insignificant by our model). The exception to this rule comes when both electricity and natural gas prices reach very high levels (90^{th} percentiles). In this case, there is a strong causality from electricity to gas that, nevertheless, disappears after approximately 8 days (with a peak on the day of the original shock above 30% in the cross-quantilogram).

All in all, we can confidently state that there is also predictive causality running from electricity to natural gas prices, and this causality seems to operate specially at the "median scenarios", in which one or the two prices are recorded near to its median values, or when high prices are induced in both markets by a relatively scarcity of the resources (that is, when both prices are recorded at the right tail of their respective distributions, 90^{th} percentiles). The causality from electricity to gas is, as expected, lower than the effect running from natural gas to electricity and it virtually disappears when at least one of the two resources are relatively cheap (10^{th} percentile). Our analysis is confirmed by the Q-statistics reported in Figure 4.

Figure 5. Cross-quantilogram to measure the directional predictability from electricity to natural gas prices: PJM

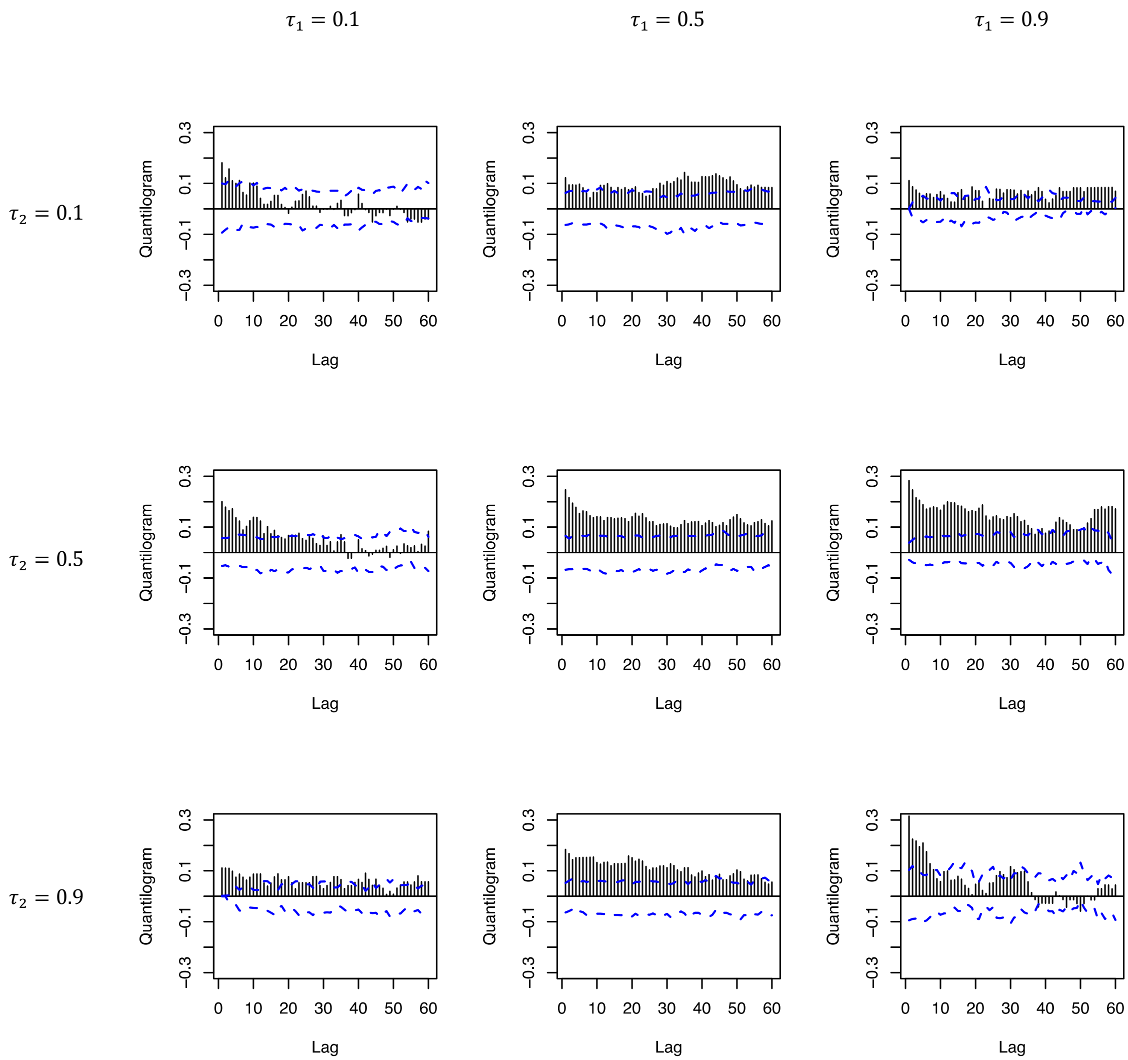

In the figure are presented the sample cross-quantilograms $\hat{\rho}_\tau(k)$ for $k = 1 \dots 60$ and $\tau_i = (0.1, 0.5, 0.9)$ for $i = 1,2$. The cross-quantilogram measures the directional predictability from electricity to natural gas prices, in PJM. In the rows are sorted the quantiles of the electricity prices, while in the columns of the natural gas prices. The bars represent $\hat{\rho}_\tau(k)$, while the dotted lines are confidence intervals at 95% of confidence constructed by the stationary bootstrap centered at zero with $B = 1000$ number of repetitions.

Figure 6. Box-Ljung test and directional predictability from electricity to natural gas: PJM

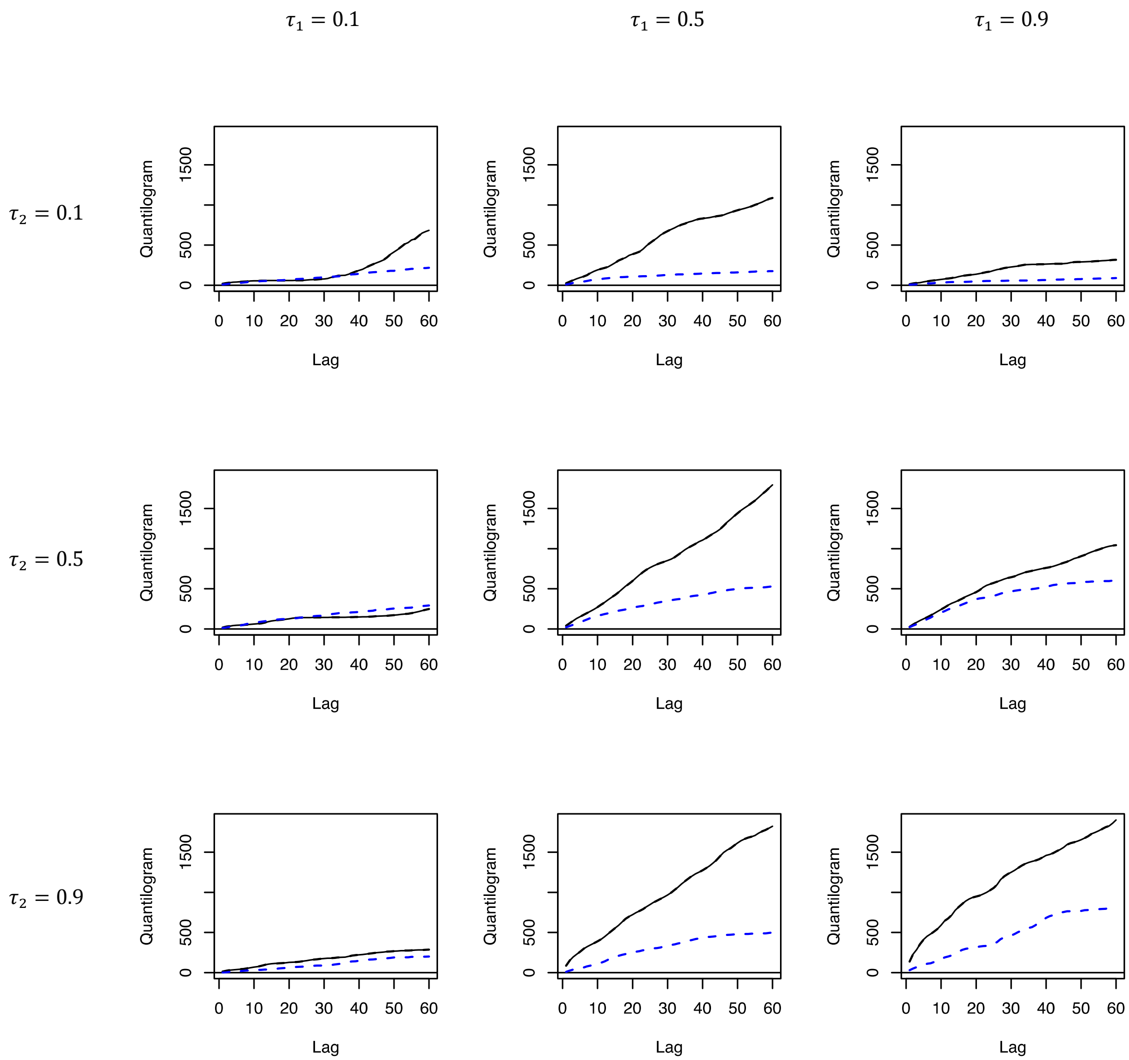

In the figure are presented the sample $\hat{Q}_{\tau}^{(p)}$ for $p = 1 \dots 60$ and $\tau_i = (0.1, 0.5, 0.6)$ for $i = 1,2$. The solid lines are the Q-Box-Ljung statistics, while the dotted lines are confidence intervals at 95% of confidence, constructed by the stationary bootstrap centered at zero with $B = 1000$ number of repetitions.

1.2. *Results for the New England interconnection*

Figure 7 depicts the cross-quantilogram with 60 lags and different quantiles, with the predictability running from natural gas to electricity prices, for the New England market. As it occurred for the PJM market, the results vary depending on the analyzed conditional quantile, both in the predicted variable (electricity power in rows) and the predicting variable (natural gas in columns). Figure 8 presents the corresponding Q-statistics that test the null of non-predictability from gas to electricity at different cumulated lags and that condition on different quantiles. The same pattern documented before arises, perhaps more clearly, when we focus on Figure 7: the higher the conditional quantiles of prices, the higher the levels of predictability. For this reason, the cross-quantilograms are larger when the figure depicts the intersection between the quantiles above the median of both electricity and natural gas prices. By the contrary when one or the two series are at their 10^{th} percentile, the predictability reduces both in magnitude and persistence. Finally, when the two prices are relatively high (90^{th} percentile) the predictability increases considerably, arising to more than 60%, either contemporaneously or with a few lags. These observations are ratified when we observe Figure 8. In none of the cases the hypothesis of predictability running from natural gas to electricity prices can be rejected, but the rejection of the null of non-predictability is stronger precisely when both prices reach their respective 90^{th} percentiles.

Figure 7. Cross-quantilogram to measure the directional predictability from natural gas to electricity prices: New England

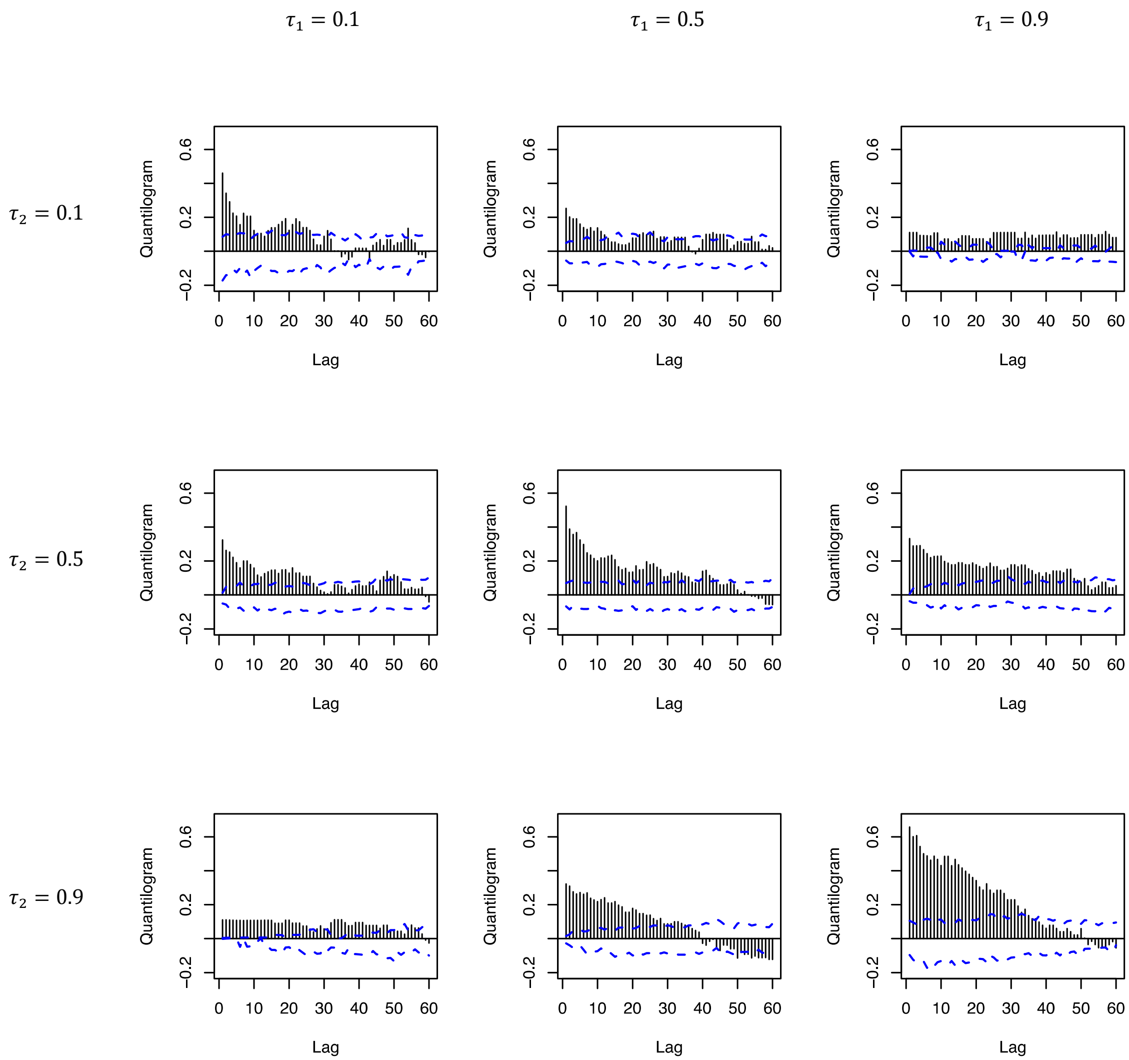

In the figure are presented the sample cross-quantilograms $\hat{\rho}_{\tau}(k)$ for $k = 1 \ldots 60$ and $\tau_i = (0.1, 0.5, 0.6)$ for $i = 1,2$. The cross-quantilogram measures the directional predictability from natural gas to electricity prices, in New England. In the rows are sorted the quantiles of the natural gas prices, while in the columns of the electricity prices. The bars represent $\hat{\rho}_{\tau}(k)$, while the dotted lines are confidence intervals at 95% of confidence constructed by the stationary bootstrap centered at zero with $B = 1000$ number of repetitions.

Figure 8. Box-Ljung test and directional predictability from natural gas to electricity prices: New England

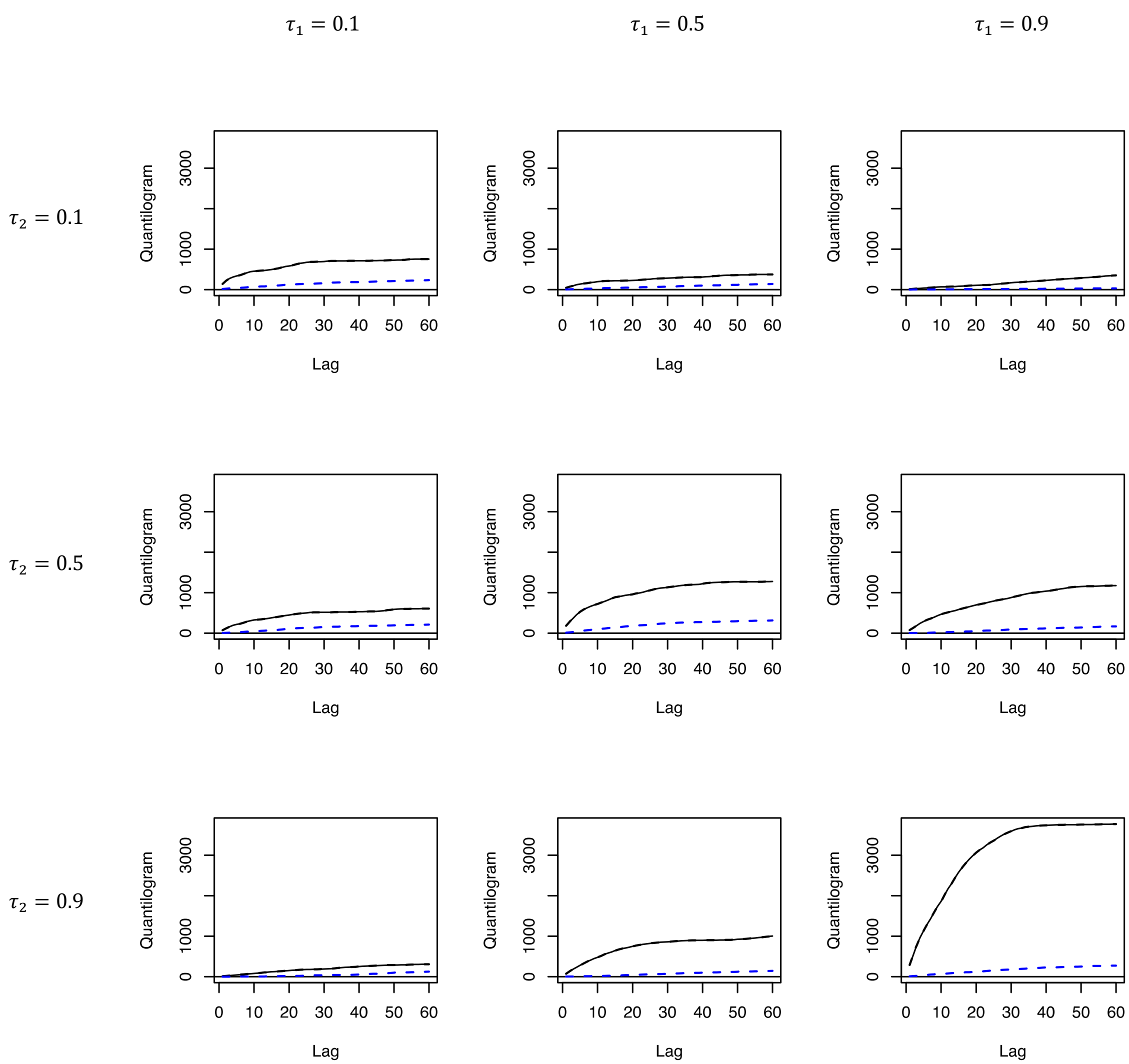

In the figure are presented the sample $\hat{Q}_{\tau}^{(p)}$ for $p = 1 \ldots 60$ and $\tau_i = (0.1, 0.5, 0.6)$ for $i = 1,2$. The solid lines are the Q-Box-Ljung statistics, while the dotted lines are confidence intervals at 95% of confidence, constructed by the stationary bootstrap centered at zero with $B = 1000$ number of repetitions.

We turn now to analyze the predictability from electricity to natural gas prices in the New England market. Figure 9 shows the estimated cross-quantilogram, with the causality running from electricity to natural gas prices, and in Figure 10 we present the respective Q-statistics. The hypothesis of the double causality cannot be rejected, this time, even clearer than for the PJM market. The magnitude of the predictability equals that observed from gas to electricity and even overpasses it at different quantiles, as for instance the median quantiles of the predicting variable (electricity). Noticeable the predictability increases considerably when we move from low to high quantiles of the predicting variable (electricity), especially when the predicted variable remains itself above its own median. In general, when one of the two prices is relatively low, that is at the 10^{th} percentile, the predictability is low as well compared to the cases of the highest quantiles in the two series. Figure 8 complements the analysis and tells us that the hypothesis of non-predictability from electricity to natural gas prices is always rejected at all quantiles in both series, but especially at the intersection of the highest quantiles.

Comparing the results for the New England interconnection with those of the PJM market, it becomes evident that predictability is larger almost always in the former market, especially when the power resources are relatively scarce and the prices are relatively high. This may be related to greater dependence in the New England market of natural-gas power generation (50% of the total), compared to the case of PJM (24%).

Figure 9. Cross-quantilogram to measure the directional predictability from electricity to natural gas prices: New England

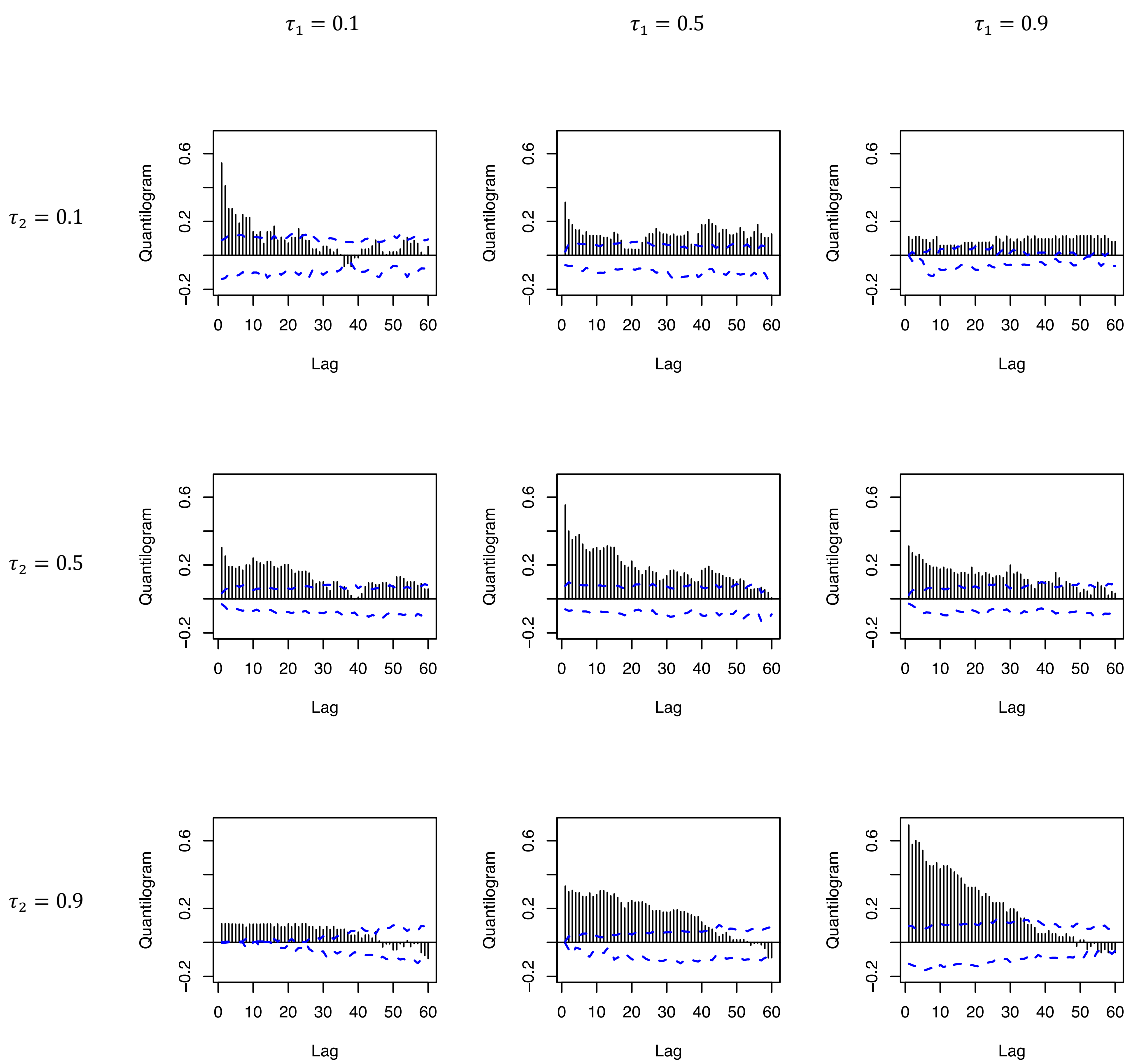

In the figure are presented the sample cross-quantilograms $\hat{\rho}_\tau(k)$ for $k = 1 \ldots 60$ and $\tau_i = (0.1, 0.5, 0.6)$ for $i = 1{,}2$. The cross-quantilogram measures the directional predictability from electricity to natural gas prices, in New England. In the rows are sorted the quantiles of the electricity prices, while in the columns of the natural gas prices. The bars represent $\hat{\rho}_\tau(k)$, while the dotted lines are confidence intervals at 95% of confidence constructed by the stationary bootstrap centered at zero with $B = 1000$ number of repetitions.

Figure 10. Box-Ljung test and directional predictability from electricity to natural gas: New England

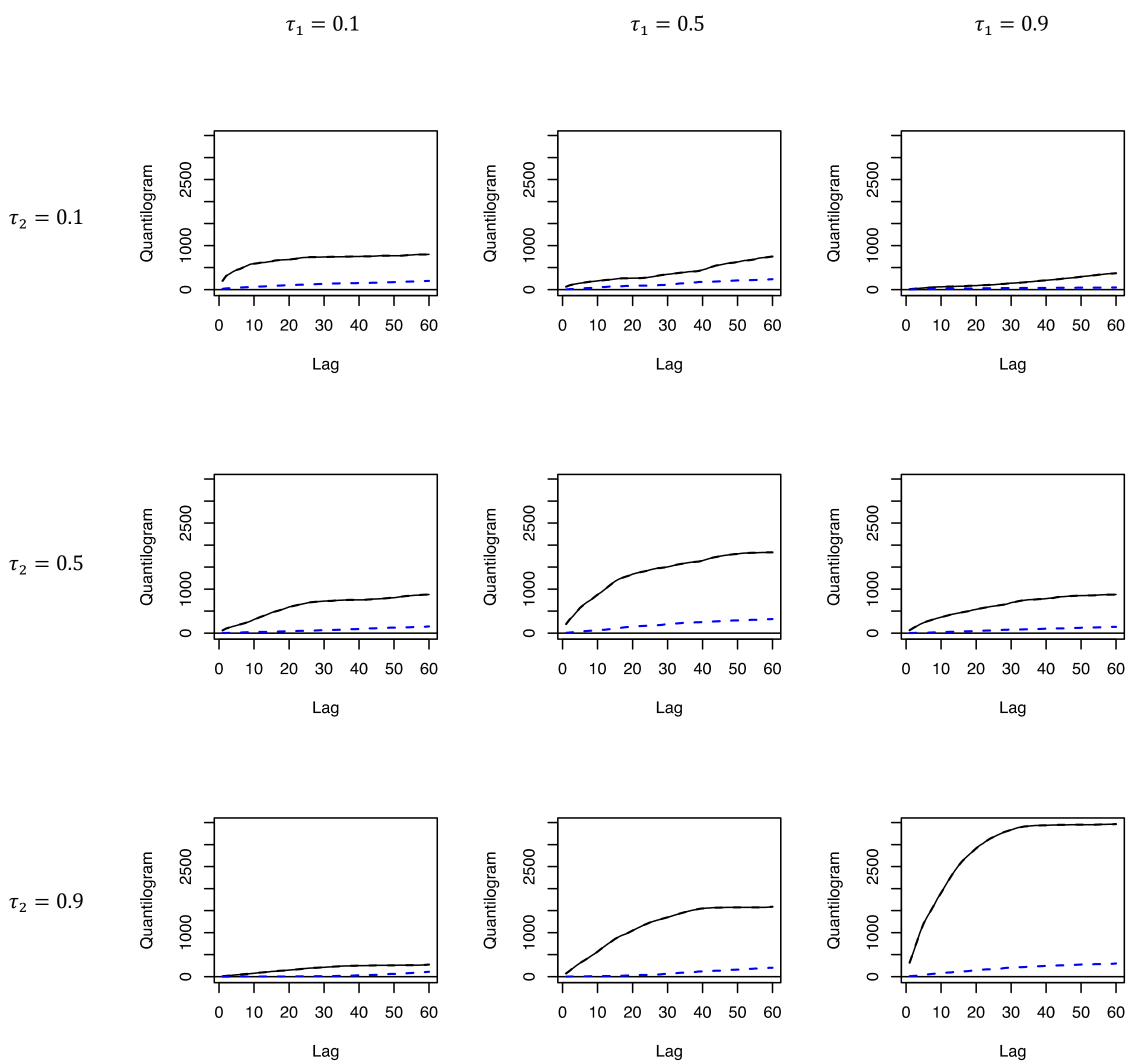

In the figure are presented the sample $\hat{Q}_{\tau}^{(p)}$ for $p = 1 \dots 60$ and $\tau_i = (0.1, 0.5, 0.6)$ for $i = 1,2$. The solid lines are the Q-Box-Ljung statistics, while the dotted lines are confidence intervals at 95% of confidence, constructed by the stationary bootstrap centered at zero with $B = 1000$ number of repetitions.

5. Conclusion

We test for directional predictability between electricity and gas prices at different quantiles of the price distribution in the PJM and the New England interconnections. The predictive causality goes in the two directions, from gas to electricity and from electricity to gas. Nevertheless, the relationship changes remarkably conditioning on different quantiles of the price distribution. In the two analyzed markets the soundest directional predictability is recorded when the electricity and gas prices are at their 90^{th} percentile. In PJM the cross-correlation between the quantile-hit functions is found to be around 20% -30% from gas to energy, and around 15% from electricity to gas, at the 90^{th} percentile of the series. In the New England market these numbers lie between 40% and 60% (in both directions) and the relation lasts in average 40 days.

In all the cases the directional predictability is considerably reduced for the lower quantiles (10^{th} percentile) of the two series, but in the case of PJM, this is especially true when the predicting variable (either gas or electricity) has reached its 10^{th} percentile. In general the cross-quantilograms are significant in the median-up ranges of the two variables, namely in the intersection between the 50^{th} and 90^{th} percentiles of the series. This informs about the existence of a significant a positive feedback effect between the two markets during both, relatively regular market scenarios, and in cases of relatively scarcity of the resources when the prices hit the 90^{th} percentile in the two markets, simultaneously. This analysis is possible by exploring the two quantile-dimensions that cross-quantilograms offer as an empirical device, and therefore are novel to the literature.

The remarkable decreasing in the price of natural gas that the world economy has witnessed during the last years (with a perspective of even further decreases), the environmental regulations targeted to force old coal based power stations to curtail their production and to phase-out of nuclear power plants, and the increment in the global supply of power coming from intermittent renewables sources, particularly wind, guarantee the future role of natural gas prices as the main determinant of electricity prices. Our analysis for the PJM and the New England interconnections shows that indeed natural gas prices can be used as a predictor of electricity prices, but also the other way around, electricity prices predict natural gas prices, especially when both resources are trading relatively high in the market (at or around the 90^{th} percentile). Regulators and policy makers must consider this growing importance of gas as a primary fuel, in order to exploit its benefits and avoid its main dangers. For example, particular attention should be placed on investing in flexible and efficient gas-based power stations, in transportation infrastructures, distribution and storage of natural gas. So as in the design of policies that foster gas power plans capacity as a factor in the main grid, and therefore may attenuate the effect of natural gas on electricity (and vice versa, given the documented feedback effect from electricity to gas), especially when energy sources are relatively scarce and therefore prices are relatively high.

The empirical evidence provided in this study argues in favor of an integrated approach when considering natural gas and electricity markets. Thus, energy policy recommendations should be bounded to an integrated approach of the two markets that considers the intertwined mechanisms that tie their prices. To reduce the over-reliance on gas, as a primary source of power, inefficient gas-fired generators should be replaced with newer facilities and renewable energy devices, which have witnessed a remarkably increase as a primary power source around the world from 2003 and on.

Our results, being conditional on the market prices, are particularly insightful about the feedback effect between natural gas and electricity prices that takes place precisely when the prices are high. These effects, and therefore the possibility of witnessing an energy crisis, increase precisely when the resources are more valued, that is, when energy is more scarce and expensive. Therefore it becomes fundamental assuring a steady and adequate supply of natural gas, especially under these circumstances, by preventing bottlenecks in natural- gas pipelines or their monopolization, and by assuring appropriated natural-gas storage capacity.

References

Alexopoulos, T. A. (2017). The growing importance of natural gas as a predictor for retail electricity prices in US. *Energy*, *137*, 219–233. http://doi.org/10.1016/j.energy.2017.07.002

Brown, S. P. A., & Yücel, M. K. (2008). Deliverability and regional pricing in U.S. natural gas markets. *Energy Economics*, *30*(5), 2441–2453. http://doi.org/10.1016/j.eneco.2007.12.015

Bunn, D., Andresen, A., Chen, D., & Westgaard, S. (2016). Analysis and forecasting of electricity price risks with quantile factor models. *The Energy Journal*, *37*(1). http://doi.org/DOI: 10.5547/01956574.37.1.dbun

Chae, Y., Kim, M., & Yoo, S. H. (2012). Does natural gas fuel price cause system marginal price, vice-versa, or neither? A causality analysis. *Energy*, *47*(1), 199–204. http://doi.org/10.1016/j.energy.2012.09.047

Hagfors, L. I., Bunn, D., Kristoffersen, E., Staver, T. T., & Westgaard, S. (2016). Modeling the UK electricity price distributions using quantile regression. *Energy*, *102*, 231–243. http://doi.org/10.1016/j.energy.2016.02.025

Han, H., Linton, O., Oka, T., & Whang, Y. J. (2016). The cross-quantilogram: Measuring quantile dependence and testing directional predictability between time series. *Journal of Econometrics*, *193*(1), 251–270. http://doi.org/10.1016/j.jeconom.2016.03.001

Koenker, R. W., & Bassett, G. (1978). Regression Quantiles. *Econometrica*, *46*(1), 33–50. http://doi.org/10.2307/1913643

Linton, O., & Whang, Y. J. (2007). The quantilogram: With an application to evaluating directional predictability. *Journal of Econometrics*, *141*(1), 250–282. http://doi.org/10.1016/j.jeconom.2007.01.004

Lobato, I. N. (2001). Testing that a dependent process is uncorrelated. *Journal of the American Statistical Association*, *96*(455), 1066–1076. http://doi.org/10.1198/016214501753208726

Mosquera-López, S., Uribe, J. M., & Manotas-Duque, D. F. (2017). Nonlinear empirical pricing in electricity markets using fundamental weather factors. *Energy*, *139*, 594–605. http://doi.org/10.1016/j.energy.2017.07.181

Nakajima, T., & Hamori, S. (2013). Testing causal relationships between wholesale electricity prices and primary energy prices. *Energy Policy*, *62*, 869–877. http://doi.org/10.1016/j.enpol.2013.07.033

Politis, D. N., & Romano, J. P. (1994). The stationary bootstrap. *Journal of the American Statistical Association*, *89*(428), 1303–1313. http://doi.org/https://doi.org/10.1080/01621459.1994.10476870

Shao, X. (2010). A self-normalized approach to confidence interval construction in time series. *Journal of the Royal Statistical Society. Series B (Statistical Methodology)*, *72*(3), 343–366.

Woo, C. K., Olson, A., Horowitz, I., & Luk, S. (2006). Bi-directional causality in California's electricity and natural-gas markets. *Energy Policy*, *34*(15), 2060–2070. http://doi.org/10.1016/j.enpol.2005.02.016