



Laboratory of Economics and Management
Sant'Anna School of Advanced Studies

Piazza Martiri della Libertà, 33 - 56127 PISA (Italy)
Tel. +39-050-883-343 Fax +39-050-883-344
Email: lem@sssup.it Web Page: <http://www.lem.sssup.it/>

LEM

Working Paper Series

An Empirically Based Model of the Supply Schedule in Day-Ahead Electricity Markets

Sandro SAPIO

LEM, Sant'Anna School of Advanced Studies, Pisa, Italy

2006/12

May 2006

ISSN (online) 2284-0400

An empirically based model of the supply schedule in day-ahead electricity markets*

Sandro Sapio[†]

July 27, 2006

Abstract

The first part of this paper establishes some new pieces of evidence on the dynamics of prices and volumes in wholesale electricity day-ahead markets (NordPool, APX, Powernext). The growth of prices is more strongly autocorrelated than the growth of volumes; it is more heavy-tailed; and its conditional standard deviation decays like the reciprocal of the price level (*1/P scaling*). In the second part of the paper, it is shown that a linear supply function with stochastic intercept and constant slope suffices to explain the *1/P scaling*. Furthermore, this model allows to decompose price fluctuations in an exogenous *demand effect* and a strategically-driven *supply effect*. In light of this model, the heavier tails of price growth and its stronger autocorrelation structure are due to persistent and intermittent strategic moves by suppliers, related to expected demand growth.

JEL Classifications: C16, D4, L94.

Keywords: Electricity Markets, Supply Curve, Subbotin Distribution, Fat Tails, Scaling, Demand Effect, Supply Effect.

1 Introduction

As outcomes of the liberalization policies pursued all around the world from the Eighties on, wholesale electricity markets challenge the economic profession, thanks to a number of characteristics (non-storability, low price-elasticity of demand, transmission constraints, network congestion) which make the optimal design of electricity auctions a complex problem and engender high opportunities for market manipulation. Most of the existing literature analyses the day-ahead market, a uniform-price, sealed-bid auction which determines equilibrium prices and quantities each day for every hour of the day after.

Stimulated by the rich structure of price fluctuations in day-ahead markets (multiple periodic patterns, persistency, spikes, heavy tails, time-dependent volatility), efforts towards proper statistical modelling of electricity price dynamics have blossomed (Geman and Roncoroni 2002, Eberlein and Stahl 2003, Weron, Bierbrauer and Truck 2004, Sapio 2004, Bottazzi,

*Preliminary draft: please, do not quote without permission. I am grateful to Valeria Termini and to NordPool SA for data availability, and to Stanley Metcalfe for very valuable comments. The usual disclaimer applies.

[†]LEM, Sant'Anna School of Advanced Studies, Pisa. E-mail: ssapio@sssup.it

Sapio and Secchi 2005, Knittel and Roberts 2005, Guerci et al. 2006, among others). Surprisingly, much less explored are the statistical properties of day-ahead power volumes dynamics:¹ though, auction-theoretic models of electricity pools (von der Fehr and Harbord, 1993) and the pivotal-generator approach (Blumsack et al., 2002) suggest that the demand-capacity ratio and its fluctuations are key in determining whether prices stay close to generating costs, or jump to the maximum admissible level. Focusing on just price dynamics can at best give a partial view of the market. Nevertheless, for any given characterization of the exogenous demand dynamics, different assumptions on the supply schedule yield different properties of price fluctuations. Insights on how day-ahead power pools work can thus be obtained by comparing distributions of both price and volume growth rates and finding those restrictions on the properties of the supply function which are consistent with the statistical evidence.

The contribution of this paper is twofold. First, it reviews and verifies some pieces of evidence on the dynamics of prices and volumes in day-ahead electricity markets, focusing on the Scandinavian *NordPool*, the Dutch *APX*, and the French *Powernext*. The comparative strength of serial correlations for the growth of prices and volumes is assessed, showing that the former are more strongly autocorrelated than the latter in all markets. Density fit exercises, based on the Subbotin family of distributions - a family including Laplace and Normal laws as special cases - show that price growth is more heavy-tailed than volume growth. Finally, the conditional standard deviation of price growth decays like the reciprocal of the price level ($1/P$ scaling) in the NordPool and Powernext markets, whereas the scaling evidence for the APX is rather mixed.

The second goal of the paper concerns the theoretical interpretation of the detected empirical facts. For the $1/P$ scaling to emerge, it suffices that the supply function is piecewise linear with stochastic intercepts and constant slopes. This model allows to decompose price fluctuations in an exogenous *demand effect* and a strategically-driven *supply effect*. Within this framework, the heavier tails of price growth and its stronger autocorrelation structure are due to persistency and intermittency in market gaming attempts by suppliers. More in detail, generating companies can respond to changing demand conditions in heterogeneous ways (e.g. entry, exit, capacity withholding) which can have opposite effects on the market price. A kurtosis gap is likely to emerge if (i) the aggregate impact of entry and exit is quantitatively greater than the impact of capacity withholding; (ii) volatile demand growth is associated with volatile shifts in the supply schedule; and (iii) individual moves are coordinated in such a way, that parallel shifts in the supply schedule are either very small, or very large. The observed gap between serial correlations of price growth and demand growth is due to persistency in the strategic responses to expected demand growth, and to adaptive behavior in the process of expectation formation.

The paper is organized as follows. In Section 2, an overview of the main empirical facts on autocorrelations (2.1), distributional shapes (2.2), and scaling (2.3) is provided. Section 3 performs a theoretical analysis of the conditions underlying the emergence of the empirical properties illustrated in Section 2, with a focus on the proposed linear supply function (3.1), and on the implied explanations of the kurtosis (3.2) and autocorrelation (3.3) patterns. Section 4 draws some conclusions and perspectives for future research.

¹See Nowicka-Zagrajek and Weron (2002) for an exception.

2 Some pieces of evidence on wholesale electricity market outcomes

The present analysis concerns three markets: NordPool (including Norway, Sweden, Finland, Denmark), 1095 daily observations from January 1, 1997, to December 31, 1999; the Amsterdam Power Exchange (APX), daily data from January 6, 2001, to December 31, 2004 (1457 observations); and the Powernext, 1065 datapoints from November 27, 2001, to December 31, 2004.² For each market, 24 prices and volumes are available for every trading day.

The structure of the markets under analysis is the following.³ Each day, 24 uniform price auctions are run simultaneously in order to determine prices and quantities for each hour of the following day. Market participants (power generating companies on the supply side, utilities and large industrial consumers on the demand side) can submit a limited number of price-quantity couples, during a given bidding period. As the bidding deadline expires, the central market operator collects all bids and asks, sorts the former in ascending order, and the latter in descending order. The 24 day-ahead market prices are determined, for each auction, by the intersection between the curves, and all power is sold and purchased at that price by inframarginal participants.

There are two ways to determine the market-clearing price. In a *flat bid order system*, the market-clearing price is determined on the basis of step curves drawn between the submitted price-volume combinations. This is true of the APX and of other markets that have not been analyzed here (such as the Spanish *Omel*, and the Slovenian *Borzen*). In the NordPool and in the Powernext, the order system is based on *interpolation* between the submitted price-volume combinations: the market supply schedule can thus be seen as a piece-wise linear function (see also Meeus, 2005). Statistical properties of market outcomes may change accordingly.

Analysis of the available datasets allows to extend and integrate the existing evidence on three phenomena: (i) autocorrelations of price growth rates at weekly lags are stronger than the corresponding autocorrelations of volume growth rates (Section 2.1); (ii) the tails of price growth rates are fatter than the tails of volume growth rates (2.2); (iii) the conditional standard deviation of price growth rates decays like the reciprocal of the price level (2.3).

2.1 Basic statistical properties

Let P_{ht} and Q_{ht} be, respectively, the price and volume at hour h of day t , and p_{ht} and q_{ht} their natural logarithms. For each given hourly market, one shall focus on the following variables:⁴

- daily price changes: $\Delta P_{ht} \equiv P_{ht} - P_{h,t-1}$
- daily price growth rates, or *log-returns*: $r_{ht} \equiv p_{ht} - p_{h,t-1} \approx \frac{\Delta P_{ht}}{P_{h,t-1}}$
- daily volume changes: $\Delta Q_{ht} \equiv Q_{ht} - Q_{h,t-1}$
- daily volume growth rates: $g_{ht} \equiv q_{ht} - q_{h,t-1}$

²Data sources, respectively: NordPool FTP Server; www.apx.nl; www.powernext.fr.

³There exists a vast literature on the institutional characteristics of electricity markets and on their evolution. Suggested overviews are in Joskow (1996), Wolak and Patrick (1997), Wilson (2000), Green (2002), Holburn and Spiller (2002), Newbery (2002), and in the book edited by Glachant and Finon (2003).

⁴Both absolute and relative changes in prices and volumes are analyzed. Indeed, while it is common to analyze log-prices, a number of authors have studied the dynamics of price levels (Alvarado and Rajaraman, 2000; Lucia and Schwartz, 2002; Knittel and Roberts, 2004).

This section gives a snapshot of the main statistical properties of the above variables, and performs comparisons between the properties of price and volume dynamics, with remarks regarding the cross-country robustness of the detected facts.

Table 1: Summary statistics of daily log-returns, price changes, load growth rates, and load changes, for the NordPool, APX, and Powernext markets: 1 am and 12 am.

Markets	mean		std.dev.		skewness		exc.kurt.	
	1 am	12 am	1 am	12 am	1 am	12 am	1 am	12 am
<i>r</i>								
NordPool	-.0006	-.0005	.078	.118	-.169	.769	11.402	13.100
APX	.0001	.0003	.284	.692	-.145	.804	16.376	3.587
Powernext	.0003	-.0001	.292	.494	.132	1.124	2.815	7.821
ΔP								
NordPool	-.0956	-.0771	6.610	13.115	-.050	.345	6.043	8.679
APX	.0016	.0103	5.032	112.217	-.136	-1.024	7.124	104.775
Powernext	.0038	-.0026	4.436	52.105	.118	2.102	2.084	208.785
<i>g</i>								
NordPool	.0003	.0004	.049	.088	.180	.768	1.419	2.272
APX	.0005	.0008	.220	.215	-.111	-.101	5.281	2.417
Powernext	.0024	.0029	.393	.413	-.071	.007	3.910	10.880
ΔQ								
NordPool	2.373	3.568	303.437	667.144	.106	.800	.865	2.896
APX	.611	.791	275.968	251.872	-.033	-.118	.654	1.043
Powernext	1.180	1.447	300.827	284.418	.229	.220	2.016	2.883

In Table 1, summary statistics for r , ΔP , g and ΔQ are provided, for two representative hours: 1 a.m. and 12 a.m. These tables show that, while drifts and asymmetries in price growth distributions are rather weak, standard deviations are clearly higher in day-time auctions than during the night, in all countries. As regards excess kurtosis, all series are leptokurtic, sometimes extremely so. The ΔP excess kurtosis is highest during the day, whereas log-returns r in APX are more leptokurtic by night. Descriptive statistics for g and ΔQ (1 a.m. and 12 a.m.) are in the lowest part of Table 1. All volume series show positive trends in mean, with different magnitudes. Daily changes in electricity demand are more volatile during the central hours of the day, except for APX. Skewness values are generally low, and as regards signs, no patterns can be detected. Excess kurtosis is always positive, with relatively low values for the variable ΔQ .

Relevant differences between indicators of price and volume dynamics appear quite clearly. Log-returns tend to be more volatile than volume growth rates.⁵ All series are leptokurtic, but log-returns and price changes are much more so, in comparison with volume changes and growth rates. All of these comparative properties are preserved even when time dependencies are filtered out, as shown in Table 2, which reports the summary statistics of price and volume fluctuations after a Cholesky filter has been applied (see Diebold, Ohanian and Berkovitz 1997 for its definition, and Bottazzi, Sapio and Secchi 2005 for a previous application).

The different statistical nature of prices and volumes is shown also by Figures 1, 2, and 3. While price series differ considerably across hours, volumes observed at different hours of the day look very similar to each other, apart from some diversities in variance

⁵A similar comparison between standard deviations of ΔP and ΔQ is not correct, in that units are different.




Figure 1: NordPool prices and volumes, from Jan 1, 1997, to Dec 31, 1999. Hours: 1 am, 12 am.

Table 2: Summary statistics of daily filtered log-returns, price changes, load growth rates, and load changes, for the NordPool, APX, and Powernext markets: 1 am and 12 am.

Markets	mean		std.dev.		skewness		exc.kurt.	
	1 am	12 am	1 am	12 am	1 am	12 am	1 am	12 am
ϵ_r								
NordPool	0.0000	0.0000	1.0000	1.0000	-0.2157	-0.3174	10.5679	7.1356
APX	0.0000	0.0000	1.0000	1.0000	-1.1935	0.7032	10.2266	2.7038
Powernext	0.0000	0.0000	1.0000	1.0000	-0.4670	0.2727	1.3977	5.9757
$\epsilon_{\Delta P}$								
NordPool	0.0000	0.0000	1.0000	1.0000	-0.0823	0.4523	5.1562	8.1691
APX	0.0000	0.0000	1.0000	1.0000	0.6009	4.9863	4.9543	63.7976
Powernext	0.0000	0.0000	1.0000	1.0000	-0.0311	6.6982	1.8630	94.7312
ϵ_g								
NordPool	0.0000	0.0000	1.0000	1.0000	0.1610	0.0531	1.9604	0.9392
APX	0.0000	0.0000	1.0000	1.0000	-0.6072	-0.3546	4.7253	1.4238
Powernext	0.0000	0.0000	1.0000	1.0000	0.3076	-1.1783	4.5467	12.6202
$\epsilon_{\Delta Q}$								
NordPool	0.0000	0.0000	1.0000	1.0000	-2.6990	0.1787	24.6974	0.8404
APX	0.0000	0.0000	1.0000	1.0000	0.0712	-0.0382	0.3942	0.5070
Powernext	0.0000	0.0000	1.0000	1.0000	0.3725	0.5540	0.9536	1.5032

(typically higher during the central hours of the day: see NordPool volumes). Although




Figure 2: APX prices and volumes, from Jan 6, 2001, to Dec 31, 2004. Hours: 1 am, 12 am.

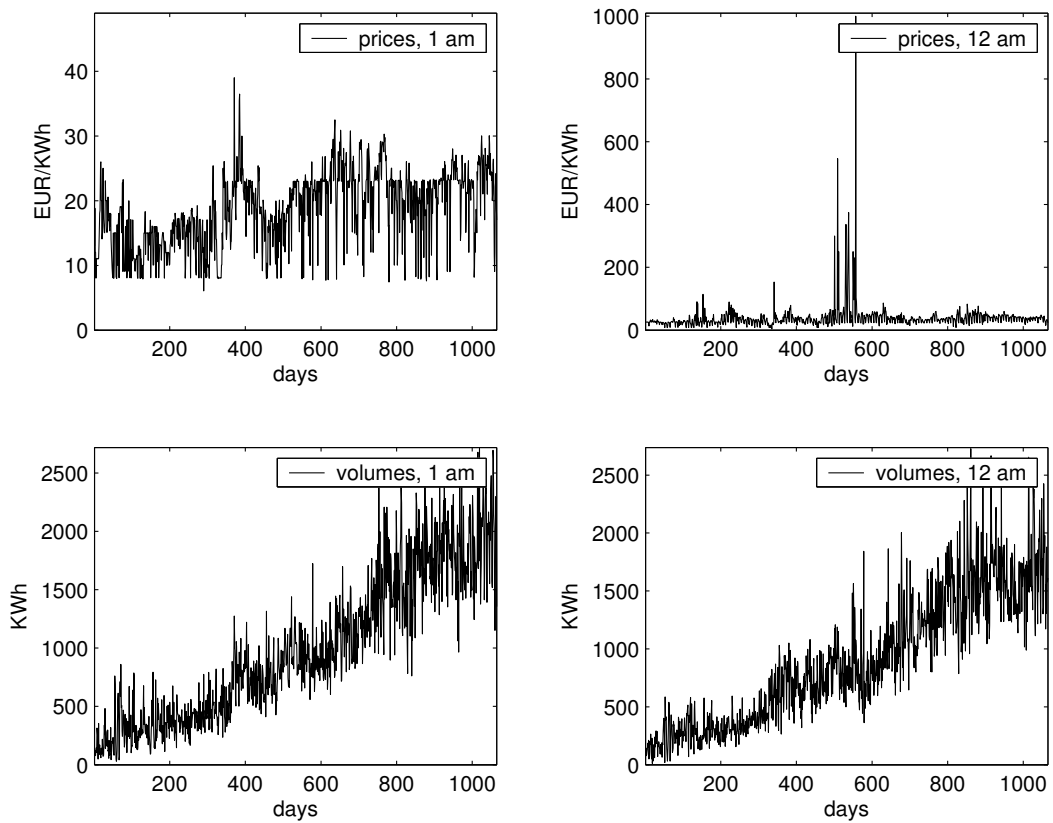


Figure 3: Powernext prices and volumes, from Nov 27, 2001, to Dec 31, 2004. Hours: 1 am, 12 am.




Figure 4: Autocorrelograms of log-returns and volume growth, 1 a.m. Autocorrelation coefficients are significant if their absolute value is greater than 0.0592 (NordPool), 0.0513 (APX), 0.0601 (Powernext).

volumes tend to display annual seasonality (NordPool, APX), this property is not always observed in prices: some annual seasonality is pretty evident in NordPool prices, much less so in APX and Powernext. In the latter market, one can moreover observe a strong increasing trend in the day-ahead volume level, which is most likely due to the increasing share of day-ahead transactions over the total national power consumption level. Finally, all price series display sharp and short-lived spikes, whereas volume series seem to be characterized by rather homoskedastic increments.

The persistent and systematic nature of time dependencies in electricity price and volume fluctuations is a well-established fact in the empirical literature on power markets (see Longstaff and Wang 2002, Weron 2002, and Sapio 2004). Along with yearly seasonals and long-run trends, the dynamics of power prices and volumes is characterized by weekly patterns. In Figures 4 and 5, this claim is corroborated by the autocorrelograms of log-returns and volume growth rates (1 a.m. and 12 a.m. series).⁶ In the 1 a.m. market, log-returns are either serially uncorrelated (NordPool) or display some very mild weekly periodicity and a small, negative autocorrelation at lag 1 day (APX, Powernext). As it can be noticed, log-returns in the 12 a.m. market are always strongly autocorrelated at lags 1 week, 2 weeks, and

⁶Autocorrelograms of ΔP and ΔQ , not reported here, are in any case very similar to those of r and g .




Figure 5: Autocorrelograms of log-returns and volume growth, 12 a.m. Autocorrelation coefficients are significant if their absolute value is greater than 0.0592 (NordPool), 0.0513 (APX), 0.0601 (Powernext).

so forth. Autocorrelation coefficients are between 0.4 and 0.6 for the first weekly lags, and decay quite slowly. Autocorrelation patterns for the NordPool volume growth mirror the ones displayed by log-returns, whereas APX and Powernext volume growth rates show first-order negative autocorrelation, and a very mild weekly pattern.

The weekly periodicity is not surprising: it is determined by decreasing energy consumption during weekends, and increasing use of power at the beginning of the week, following the overall economic activity. These patterns cannot be smoothed out, because electricity is not storable. Less obviously, the autocorrelation coefficient of price growth at lag 7 days tends to be at least as high as the corresponding autocorrelation of volume growth rates. Such an evidence is interesting: given that electricity demand is exogenous, inelastic, and tied to the time profile of the overall economic activity, serial correlations of electricity prices are expected to track demand fluctuations quite closely. The stronger structure of price dynamics may thus be due to some endogenous mechanism, such as strategic interaction between power market participants.

2.2 Subbotin density fit

A second piece of evidence on wholesale electricity markets concerns the shape of the probability density functions of indicators of price and volume dynamics. Density fit exercises based on the Subbotin family allow to quantify their degrees of peakedness and heavy-tailedness within a quite general, parsimonious, and flexible framework.

First used in economics by Bottazzi and Secchi (2003), the Subbotin probability density function of a generic random variable X reads (Subbotin, 1923):

$$Pr \{X = x\} = \frac{1}{2ab^{1/b}\Gamma(1 + \frac{1}{b})} e^{-\frac{1}{b}|\frac{x-\mu}{a}|^b} \quad (1)$$

with parameters a (width), b (shape), and μ (position). $\Gamma(\cdot)$ is the gamma function. The Subbotin reduces to a Laplace if $b = 1$ and to a Gaussian if $b = 2$.⁷ The Continuous Uniform is a limit case for $b = \infty$. As b gets smaller, the density becomes heavier-tailed and more sharply peaked.⁸

Compared to previously fitted distributions, such as the Generalized Hyperbolic (Eberlein and Stahl, 2003), the Subbotin has a more parsimonious specification: just 3 parameters need to be estimated (2 if the data are demeaned). Moreover, the Subbotin distribution allows for greater flexibility with respect to the tail behavior. For instance, the Generalized Hyperbolic distribution family only admits exponential tail decay (cf. the application by Eberlein and Stahl, 2003, on NordPool data). On the other hand, evidence of Subbotin distributions would rule out power-law tails, which characterize Levy phenomena with tail index $\alpha < 2$ (see Bystroem, 2001, Bellini, 2002, and Deng, Jiang and Xia, 2002, for applications to electricity price dynamics).

Estimates of the Subbotin parameters have been obtained through Maximum Likelihood methods, making use of the Subbotools (details are available in Bottazzi, 2004). Tables 3 and 4 display, for each of the markets under analysis, the estimated shape parameter b and the corresponding standard error, respectively for price and volume dynamics. For the sake of brevity, estimated a 's have been omitted. Furtherly, normalization implies $\mu = 0$. Figures 6, 7 and 8 depict histograms of the variables under analysis for some hours, with superimposed Subbotin fits.

Filtered log-returns display a Laplace shape, a fact holding robustly across markets, with some slight deviation (e.g. b values clustering around 0.8 during night hours in the APX). Values of b for filtered price changes are clustered around 1 for the NordPool, around 0.6-0.7 for APX and Powernext (night auctions), and higher than 1 for APX and Powernext (day auctions). Estimates of b for filtered volume growth in NordPool and APX tend to be half-way between the values implied by Laplace and Normal laws, i.e. $b \approx 1.5$. On the contrary, in the Powernext b values are around 1.1. Estimated shape coefficients for volume changes are around 1.5 or 1.6 in the NordPool, 1.7 or 1.8 in the APX, and 1.4 in the Powernext.

These results shed light on a very interesting property. Estimated shape coefficients for price dynamics variables are systematically higher than for volume dynamics, implying that

⁷The Laplace distribution (also known as *double exponential*) has been detected in diverse economic phenomena: from price changes in financial markets (Kozubowski and Podgorski, 2001; McCauley, 2004) to the growth of economic organizations (firms in Stanley et al., 1996, and in Bottazzi and Secchi, 2003; countries in Lee et al., 1998).

⁸It has been proved by West (1987) that a Subbotin with $b \in [1, 2]$ can be obtained as a mixture of normal distributions, with α -stable mixing density. Specifically, $b = 2\alpha$, showing that, when $\alpha = 0.5$, a Laplace law emerges through an Exponential mixing density.

Table 3: The estimated Subbotin shape parameters of the daily filtered log-returns and price changes for the NordPool, APX, and Powernext markets.

	r						ΔP					
	NordPool		APX		Powernext		NordPool		APX		NordPool	
	b	$std.err.$	b	$std.err.$	b	$std.err.$	b	$std.err.$	b	$std.err.$	b	$std.err.$
1	.967	.054	1.081	.054	1.259	.076	1.123	.065	1.128	.056	1.353	.083
2	.899	.049	.852	.040	1.224	.073	1.122	.065	1.255	.064	1.464	.091
3	.901	.049	.864	.041	1.222	.073	1.020	.057	1.312	.068	1.495	.094
4	.911	.050	.837	.039	1.187	.070	1.123	.065	1.380	.073	1.542	.098
5	.927	.051	.806	.038	1.056	.061	1.155	.067	1.347	.070	1.605	.103
6	.885	.048	.785	.036	1.086	.063	1.152	.067	1.318	.068	1.603	.103
7	.930	.051	.754	.035	.948	.053	1.142	.066	1.306	.068	1.391	.086
8	.944	.052	.778	.036	1.040	.060	.991	.056	.984	.048	1.280	.077
9	.956	.053	.846	.040	.998	.057	.816	.044	.772	.036	.920	.051
10	.995	.056	.899	.043	.942	.053	.884	.048	.709	.032	.648	.034
11	1.146	.066	1.023	.050	1.082	.063	1.052	.060	.685	.031	.634	.033
12	1.146	.066	1.233	.063	1.022	.058	1.140	.066	.701	.032	.641	.034
13	1.047	.059	1.121	.056	1.050	.060	1.208	.071	.652	.029	.628	.033
14	1.089	.062	1.153	.058	1.073	.062	1.196	.070	.671	.030	.640	.033
15	1.059	.060	1.103	.055	1.042	.060	1.140	.066	.670	.030	.631	.033
16	1.072	.061	1.026	.050	1.059	.061	1.018	.057	.591	.026	.624	.033
17	1.047	.059	1.045	.051	1.220	.073	.880	.048	.667	.030	1.195	.071
18	.986	.055	1.068	.053	1.147	.067	.790	.042	.670	.030	1.033	.059
19	1.069	.061	.976	.047	1.021	.058	1.020	.058	.640	.029	.967	.055
20	1.005	.056	1.102	.054	1.027	.059	1.101	.063	.714	.032	.993	.056
21	.936	.052	1.084	.054	1.149	.068	1.087	.062	.785	.036	1.264	.076
22	.972	.054	.811	.038	1.125	.066	1.126	.065	1.007	.049	1.325	.081
23	.998	.056	.772	.036	1.130	.066	1.185	.069	.987	.048	1.411	.087
24	1.029	.058	.864	.041	1.063	.061	1.152	.067	.974	.047	1.345	.082

Table 4: The estimated Subbotin shape parameters of the daily filtered volume growth and volume changes for the NordPool, APX, and Powernext markets.

	g						ΔQ					
	NordPool		APX		Powernext		NordPool		APX		NordPool	
	b	$std.err.$	b	$std.err.$	b	$std.err.$	b	$std.err.$	b	$std.err.$	b	$std.err.$
1	1.421	.087	1.408	.074	1.038	.060	.383	.018	1.715	.096	1.492	.094
2	1.424	.087	1.490	.080	1.108	.065	1.556	.097	1.811	.103	1.358	.083
3	1.097	.063	1.364	.071	1.101	.064	1.641	.104	1.723	.096	1.413	.087
4	1.436	.088	1.622	.089	1.197	.071	1.042	.059	1.884	.108	1.322	.080
5	1.435	.088	1.641	.090	1.231	.074	1.601	.101	1.889	.108	1.419	.088
6	1.398	.085	1.541	.083	1.100	.064	1.589	.100	1.800	.102	1.404	.087
7	1.348	.081	1.482	.079	1.113	.065	1.593	.100	1.808	.102	1.446	.090
8	1.514	.094	1.429	.076	1.252	.075	1.495	.093	1.737	.097	1.214	.072
9	1.483	.092	1.375	.072	1.150	.068	1.573	.099	1.561	.085	1.292	.078
10	1.519	.095	1.558	.085	1.079	.062	1.548	.097	1.626	.089	1.377	.085
11	1.487	.092	1.366	.072	1.169	.069	1.557	.098	1.546	.084	1.391	.086
12	1.467	.090	1.398	.074	1.164	.069	1.569	.099	1.734	.097	1.391	.086
13	1.516	.094	1.479	.079	1.151	.068	1.551	.097	1.724	.096	1.381	.085
14	1.484	.092	1.531	.083	1.118	.065	1.631	.104	1.693	.094	1.313	.080
15	1.540	.096	1.539	.083	1.119	.065	1.667	.107	1.713	.096	1.275	.077
16	1.570	.099	1.499	.081	1.155	.068	1.614	.102	1.708	.095	1.303	.079
17	1.530	.095	1.470	.079	1.025	.059	1.577	.099	1.785	.101	1.456	.091
18	1.578	.099	1.624	.089	.998	.057	1.590	.100	1.927	.111	1.543	.098
19	1.610	.102	1.374	.072	1.076	.062	1.677	.107	1.679	.093	1.600	.102
20	1.607	.102	1.466	.078	1.087	.063	1.697	.109	1.851	.106	1.493	.094
21	1.494	.093	1.435	.076	1.148	.067	1.779	.116	1.875	.107	1.460	.091
22	1.346	.081	1.399	.074	1.087	.063	1.668	.107	1.774	.100	1.521	.096
23	1.463	.090	1.403	.074	1.120	.065	1.513	.094	1.782	.101	1.404	.087
24	1.498	.093	1.454	.077	1.102	.064	1.544	.097	1.814	.103	1.465	.092




Figure 6: Subbotin density fit for NordPool filtered log-returns, 12 a.m. (left), and filtered load changes, 12 a.m. (right). Notice the log scale on vertical axes.

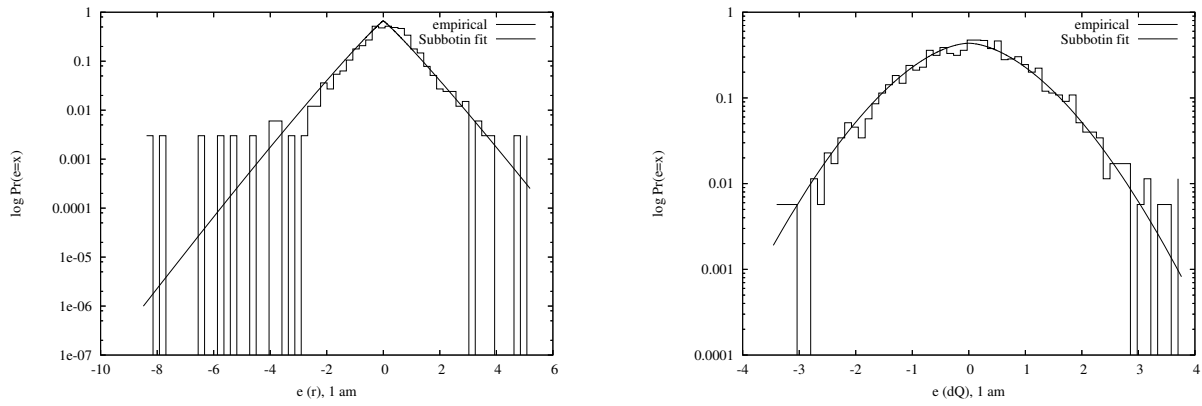


Figure 7: Subbotin density fit for APX filtered log-returns, 1 a.m. (left), filtered volume growth, 1 a.m. (right). Notice the log scale on vertical axes.




Figure 8: Subbotin density fit for Powernext filtered log-returns, 12 a.m. (left), and filtered volume changes, 1 a.m. (right). Notice the log scale on vertical axes.

the r and ΔP densities are characterized by heavier tails than for g and ΔQ . As a conclusion, extremely large relative fluctuations are more likely to be observed in electricity prices than in volumes.

2.3 Volatility-price scaling

A third fact, detected by Bottazzi, Sapio and Secchi (2005) and by Simonsen (2005) for the NordPool market, is that the conditional standard deviation of price growth rates is price-dependent.

As in previous works (Bottazzi, Sapio and Secchi, 2005; Sapio, 2005), the price growth conditional standard deviation is modelled as a power function of the price level. Taking natural logarithms, this reads

$$\log \sqrt{V[r_t|P_{t-1}]} = \chi + \rho \log P_{t-1} + \epsilon_t \quad (2)$$

where χ and ρ are constant coefficients, and ϵ_t is an i.i.d. error term. The parameter ρ is null under a multiplicative random walk; $|\rho| \in (0, 1)$ for multiplicative stationary ARMA processes with non-zero first-order autoregression; and $\rho = -1$ for all additive processes (stationary or not) and for multiplicative stationary processes with null first-order autoregression.⁹ This approach is an alternative to the common GARCH approach, which has been often applied to electricity prices (Bellini, 2002; Worthington et al., 2002; Knittel and Roberts, 2004; Cavallo, Sapio and Termini, 2005; Guerci et al., 2006).

Estimation of the power law scaling model is implemented as follows. For any given time series, data are grouped into equipopulate bins. Next, sample standard deviations of log-returns in each bin are computed, and the logarithm of the sample standard deviations is regressed on a constant and on the logarithm of the median price level within the corresponding bins.¹⁰

Table 5 reports the estimated slope coefficients ρ from Eq. 2, along with confidence intervals and R^2 values.¹¹ Three markets (NordPool, APX, and Powernext) and 24 daily series are analyzed. As a dependent variable in the regressions, the sample standard deviation of *normalized* log-returns is considered. Information about the estimated intercept is not reported, for the sake of brevity.

Point estimates of ρ suggest that standard deviations of normalized log-returns are negatively related with lagged price levels : $\rho < 0$ in almost all cases. More in detail, ρ in the NordPool and in the Powernext is often around -1, quite homogeneously across hours. It is instead difficult to detect patterns in the APX values of ρ : values around 0 are detected between 7 and 9 p.m., around -0.5 in the early afternoon, and roughly equal to -1 in all other hours. The reported bounds of the confidence intervals are very narrow, and R^2 values are generally high. Fig. 9 depicts some examples of the estimated scaling relationships.

The foregoing analysis supports evidence of a negative volatility-price relationship, at least for NordPool and Powernext:

⁹However, in multiplicative ARMA models, χ is not independent of P_{t-k} , $k = 2, 3, \dots$

¹⁰Alternatively, one could choose the mean prices within bins, or also minima or maxima. The median is a better choice because it is less affected by extreme values.

¹¹Each time series has been split into 8 equipopulate bins. The regression is run excluding the bin corresponding to the highest price levels: the latter are affected by short-lived spikes, and as such they cannot be considered homogeneous to all other price bins. Results are robust with respect to the number of bins considered.

$$\sqrt{V[r_t|P_{t-1}]} \sim \frac{1}{P_{t-1}} \quad (3)$$

The detected scaling relationship is dubbed $1/P$ scaling. The scaling coefficient for the APX market is rather unstable. Notice, in passing, that while the former markets have adopted an interpolation order system, in the latter the equilibrium price is set according to a flat bid system. Differences in statistical properties may therefore be due to differences in market settings.

3 An empirically based model of the supply curve

A broadly unexplored perspective in empirical studies of day-ahead electricity markets acknowledges that, given an inelastic demand schedule and a uniform-price auction format, one can derive the properties of price growth through the statistical regularities observed in power demand growth, and that such properties vary under different assumptions about the supply curve. Intuitively, only some supply curve models are consistent with the existing joint evidence on prices and volumes. Aim of this section is to provide some theoretical hints at the determinants of the empirical facts uncovered in this paper, namely:

1. the standard deviation of price growth goes like the reciprocal of the price ($1/P$ scaling);
2. the kurtosis of price changes is higher than the kurtosis of changes in demand;
3. the autocorrelation coefficients of price changes are higher than for changes in demand.

Let P_t and Q_t denote the wholesale price and supply of electricity at a given hour of day t , and assume they are random variables with values in \mathfrak{R}_+ . Let the inverse supply function $f : M \subset \mathfrak{R}_+ \rightarrow \mathfrak{R}_+$ be continuous and differentiable:

$$P_t = f(Q_t; \theta_t) \quad (4)$$

where M is a neighbourhood of the market-clearing price-quantity couple, and θ_t is a set of (possibly random) parameters. This can be seen as a reduced form of a two equation system, comprising (i) an upward-sloping inverse supply function, and (ii) a stochastically-shifting inelastic demand. The non-storability of electrical power imposes the equality between supply and demand, therefore Q_t here is also interpreted as the market demand level.

Define $\Delta P_t \equiv P_t - P_{t-1}$ and $\Delta Q_t \equiv Q_t - Q_{t-1}$ as respectively the change in price and demand from $t - 1$ to t . Taking the Taylor approximation of P_t about Q_{t-1} and subtracting P_{t-1} yields

$$\Delta P_t \approx \Delta f_t(Q_{t-1}) + f'(Q_{t-1}; \theta_t) \Delta Q_t + \frac{1}{2} f''(Q_{t-1}; \theta_t) (\Delta Q_t)^2 \quad (5)$$

where $\Delta f_t(Q_{t-1}) \equiv f(Q_{t-1}; \theta_t) - f(Q_{t-1}; \theta_{t-1})$ - i.e. the shift of the function $f(\cdot)$ evaluated at Q_{t-1} . This formulation allows to explicitly link the dynamics of the electricity price with the fluctuations of volumes.

The term $\Delta f_t(Q_{t-1})$ can be seen as the component of electricity price fluctuations which is not directly due to shifts in demand (it is the value of ΔP_t when $\Delta Q_t = 0$). Such a component conveys the effects of all supply-side phenomena which affect the wholesale price of electricity

Table 5: Slopes of the power law relationship between standard deviation of normalized daily log-returns and lagged price levels. Markets: NordPool, APX, and Powernext. Number of equipopulate bins: 8.

Hour	NordPool		APX		Powernext	
	ρ	R^2	ρ	R^2	ρ	R^2
1	$-1.183 \pm .021$.738	$-1.194 \pm .014$.860	$-.570 \pm .010$.751
2	$-1.218 \pm .022$.715	$-1.151 \pm .029$.580	$-.778 \pm .015$.702
3	$-1.182 \pm .018$.793	$-.987 \pm .015$.798	$-.783 \pm .012$.801
4	$-1.191 \pm .014$.869	$-.744 \pm .019$.576	$-1.065 \pm .016$.797
5	$-1.128 \pm .012$.892	$-.618 \pm .013$.676	$-1.065 \pm .016$.797
6	$-1.260 \pm .013$.885	$-.747 \pm .008$.890	$-1.275 \pm .016$.855
7	$-1.196 \pm .016$.834	$-.459 \pm .012$.549	$-1.275 \pm .016$.855
8	$-1.102 \pm .008$.938	$-.450 \pm .014$.468	$-.978 \pm .013$.841
9	$-.958 \pm .019$.678	$-.852 \pm .016$.712	$-.978 \pm .013$.841
10	$-1.034 \pm .022$.668	$-.996 \pm .013$.837	$-1.232 \pm .015$.863
11	$-1.121 \pm .014$.840	$-.988 \pm .016$.758	$-1.232 \pm .015$.863
12	$-1.298 \pm .011$.923	$-.552 \pm .011$.694	$-1.005 \pm .020$.696
13	$-1.431 \pm .008$.966	$-.625 \pm .010$.772	$-1.299 \pm .016$.856
14	$-1.313 \pm .008$.957	$-.505 \pm .013$.569	$-1.376 \pm .017$.849
15	$-1.308 \pm .006$.974	$-.528 \pm .013$.583	$-1.298 \pm .013$.896
16	$-1.283 \pm .008$.962	$-.654 \pm .013$.673	$-1.165 \pm .017$.808
17	$-1.071 \pm .012$.868	$-.792 \pm .006$.936	$-1.025 \pm .015$.827
18	$-.943 \pm .017$.732	$.016 \pm .021$.001	$-.876 \pm .005$.957
19	$-1.081 \pm .014$.834	$.595 \pm .020$.424	$-1.124 \pm .009$.929
20	$-1.244 \pm .012$.909	$.377 \pm .010$.558	$-1.124 \pm .009$.929
21	$-1.452 \pm .009$.957	$.332 \pm .019$.216	$-1.164 \pm .013$.879
22	$-1.459 \pm .007$.971	$-.886 \pm .066$.134	$-1.514 \pm .009$.960
23	$-1.368 \pm .014$.894	$-.264 \pm .059$.017	$-1.556 \pm .005$.986
24	$-1.206 \pm .010$.923	$-1.523 \pm .060$.359	$-1.665 \pm .021$.841

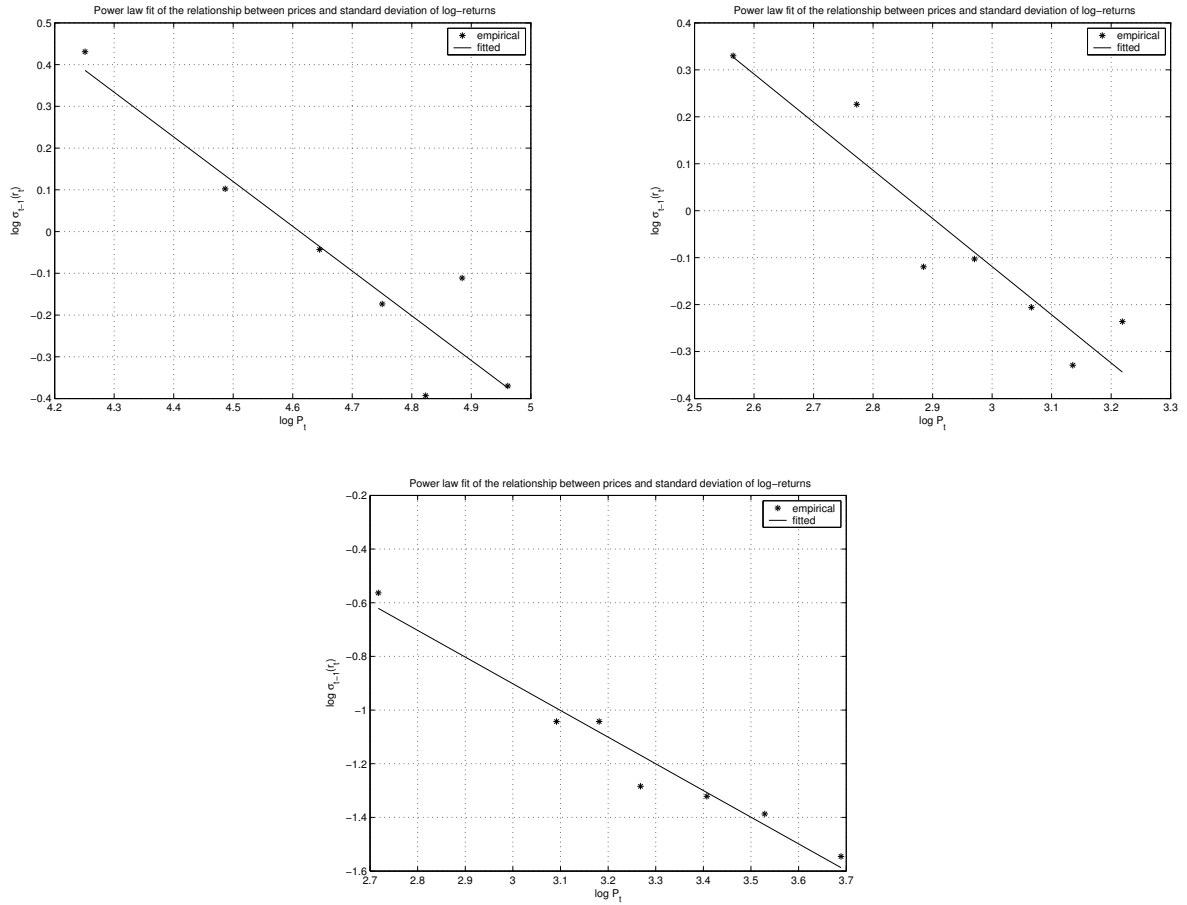


Figure 9: Linear fit of the relationship between log of the conditional standard deviation of normalized log-returns, $\log \sigma_{t-1}(r_t)$, and lagged log-price level $\log(P_{t-1})$. Clockwise: NordPool (left, 5 p.m., $\beta = -1.071$), APX (left, 1 a.m., $\beta = -1.194$), Powernext (7 p.m., $\beta = -1.124$).

between $t - 1$ and t - such as plant failures, strategic capacity withholding, changes in the level of water reservoirs, the dynamics of fuel costs, changes in markups - holding constant the level of electricity demand at $t - 1$. Henceforth, it will be referred to as *supply effect*. The remaining component is a polynomial function of ΔQ_t : it measures the direct impact of changed demand levels on the electricity price, holding constant the supply schedule. Therefore it is called *demand effect*. Yet, the impact of demand is tuned by parameters (f' , f'' , and so forth) which may reflect the strategic interaction between agents on the market - i.e. supply-side phenomena. More generally, the shape of the supply function, as defined by the parameters Δf , f' , f'' etc., is supposed to depend on the expectations formed by market participants as regards the magnitude and sign of demand growth. Thus, the supply-effect and demand-effect components may not be independent. This sort of *electricity price growth accounting* can however be extremely useful in guiding the economic interpretation of the statistical evidence presented in this paper.

The first part of this section (3.1) shows that, under certain conditions, higher order terms of the Taylor approximation are not needed for explaining the $1/P$ scaling. A linear model with constant slope and (possibly random) intercept suffices. Starting from this simple model, the second part of the section (3.2) restricts the set of possible combinations of parameter values and properties which is consistent with the observed patterns of kurtosis in the distributions of price and demand changes. In the third and final part (3.3), autocorrelation in intercept shifts is shown to strengthen the patterns of serial dependence of demand changes, yielding an even stronger autocorrelation structure in price dynamics.

3.1 Volatility-price scaling and linear supply

The empirical evidence on wholesale electricity markets reveals that the conditional standard deviation of price growth goes like the reciprocal of the price level. Sufficient conditions for the model of Eq. 5 to reproduce the scaling evidence are proposed as follows.

Proposition. Consider the model

$$r_t \approx \frac{1}{P_{t-1}} [\Delta f_t(Q_{t-1}) + f'(Q_{t-1}; \theta_t) \Delta Q_t + \frac{1}{2} f''(Q_{t-1}; \theta_t) (\Delta Q_t)^2] \quad (6)$$

The conditional standard deviation of price growth rates goes like the reciprocal of the price level under the following sufficient conditions:

(i) Suppose $V[(\Delta Q_t)^k | Q_{t-1}]$ is independent of Q_{t-1} and P_{t-1} , $\forall k = 1, 2, \dots$. Then the sufficient condition is that the inverse supply function f is linear, with (possibly) random intercept and constant slope,

$$P_t = \alpha_t + \beta Q_t \quad (7)$$

and that both slope and intercept are independent of past values of price and demand.

(ii) Suppose $V[(\Delta Q_t)^k | Q_{t-1}]$ is *not* independent of Q_{t-1} and P_{t-1} . Then the sufficient condition is that $V[\Delta f_t]$ and $f^{(k)}(Q_{t-1})(\Delta Q_t)^k$ are independent of Q_{t-1} and P_{t-1} , for $k = 1, 2, \dots$

Proof. It suffices to show that the variance of ΔP_t is independent of P_{t-1} and of Q_{t-1} . Independence from P_{t-1} is not enough, because $P_{t-1} = f(Q_{t-1}; \theta_{t-1})$. Two cases are considered.

(i) If f' is constant with respect to Q_{t-1} , all higher derivatives of f are zero. Moreover, Δf_t is independent of volumes and prices at $t - 1$. The supply schedule is therefore locally linear. The variance of r_t in such a case reads:

$$V[r_t|P_{t-1}] = \frac{V[\Delta f_t + f'\Delta Q_t]}{P_{t-1}^2} \quad (8)$$

Given $V[\Delta Q_t|Q_{t-1}]$ independent of prices and volumes at $t - 1$, it is easy to see that if f' and the conditional variance of Δf are independent as well, the variance of ΔP_t is independent, too. Hence, $1/P$ scaling results. Notice that the scaling evidence can be reproduced even if $V[\Delta f_t] = 0$, i.e. a constant intercept. The slope β must also be constant over time. This is required for $V[\Delta f_t]$ to be independent of Q_{t-1} . Indeed, if $\beta_t \neq \beta_{t-1}$, then $\Delta f_t(Q_{t-1}) = \Delta\alpha_t + \Delta\beta_t Q_{t-1}$, which conditional variance depends on Q_{t-1} , and thus on P_{t-1} .

(ii) The variance of r_t in such a case reads:

$$V[r_t|P_{t-1}] = \frac{V[\Delta f_t + f'\Delta Q_t + f''(\Delta Q_t)^2]}{P_{t-1}^2} \quad (9)$$

Suppose $V[\Delta Q_t|Q_{t-1}]$ is not independent of prices and volumes at $t - 1$. Then it can be seen that the conditional variance of ΔP_t (the above numerator) is independent as well if the conditional variance of Δf_t is independent, too, and if f', f'' etc. are not independent of prices and volumes in the previous market session, in such a way that variances $V[f^{(k)}(\Delta Q_t)^k]$ are independent.

Case (i) is consistent with a broad family of volume generating processes, such as any additive time series (either stationary or not), and multiplicative models with null AR coefficient at lag one. The latter may look too restrictive. Though, even if positive and high autocorrelation is usually observed at lag 1 day, this may in fact reflect annual - rather than 1st order - autocorrelation. The remainder of the paper shall deal with case (i), while leaving to future research the investigation of case (ii).

In terms of the supply effect-demand effect dichotomy, the above proposition suggests that the $1/P$ scaling relationship can be obtained by assuming a time-varying supply effect, and a demand effect characterized by a constant factor of proportionality. In other words, the exogenous dynamics of demand and the endogenous strategic response by market participants are additive. All effects of market gaming result in intercept shifts, without affecting the slope of the supply schedule at the intersection between demand and supply.

For a further illustration of the above proposition, let us assess the predictions of an existing supply function model.

Mount's model. Mount's (2000) analysis aims to understand how randomness about the electricity load is amplified by the structure of offers to sell power, and to explain why, even if the load is symmetrically distributed, some positive skewness can emerge in the distribution of prices. Mount put forward a mixed linear-power model:

$$P_t = \alpha + Q_t^\beta \quad (10)$$

Positive price skewness can be predicted by the above equation if one assumes $\beta > 1$, i.e. convexity of the market supply function - an assumption made plausible by convexity of short-run variable costs of power generation. However, Mount did not derive any implication as to the distribution of price changes. Mount's model is in fact ruled out by the evidence on

volatility-price scaling. To see this, let $f(Q_t; \alpha, \beta)$ be the inverse supply function, and let us compute $\Delta f_t(Q_{t-1})$ and $f'(Q_{t-1})$:

$$\Delta f_t(Q_{t-1}) = Q_{t-1}^\beta - Q_{t-1}^\beta = 0 \quad (11)$$

i.e. a null supply effect; but

$$f'(Q_{t-1}) = \beta Q_{t-1}^{\beta-1} \quad (12)$$

Clearly, in this case the slope of f evaluated in Q_{t-1} is not independent of Q_{t-1} , so that it depends indirectly upon P_{t-1} . If $V[\Delta Q_t | Q_{t-1}] \sim Q_{t-1}^{2(1-\beta)}$, then the model predicts $1/P$ scaling. If instead the variance of price changes is independent of Q_{t-1} (case (i) in the Proposition), Mount's model runs against the evidence of $1/P$ volatility-price scaling.

3.1.1 Alternative interpretations of linearity

A linear description of the inverse supply function in a wholesale electricity market can sound simplistic - particularly the prediction of a constant slope - and is subject to at least two routes of criticism.

First, if rated and declared plant capacities, variable costs of power generation, and individual offer strategies are heterogeneous, why should a linear supply function emerge? Second, power generating companies can post a limited number of price-quantity couples: this should engender a step-wise supply schedule, rather than a linear, upward-sloping one.

The former doubt is justified if one views the supply schedule as globally linear. This amounts to assuming that the slope of the supply function is constant $\forall Q \in [0, \infty)$. Such an interpretation is hardly acceptable, but not really necessary. Demand fluctuations are never so wide as to cover the whole range of feasible capacity values. One can think of demand fluctuations being bounded within a certain value region, and therefore that linearity holds (at least approximately) within that region, without necessarily holding everywhere.

The second criticism, based on the step-wise nature of supply schedules, implicitly assumes that a *flat bid order system* is at work - namely, that the system marginal price is determined on the basis of step curves drawn between the submitted price-volume combinations. This is true of the APX and of other markets that have not been analyzed here (such as the Spanish *Omel*, and the Slovenian *Borzen*). In the NordPool and in the Powernext, the order system is based on *interpolation* between the submitted price-volume combinations: the market supply schedule can thus be seen as a piece-wise linear increasing function (see also Meeus, 2005).¹²

These considerations hint at some reasons why the slope of the supply function, at the market-clearing intersection, may appear constant. A first possibility is that, as demand changes, different generating companies become marginal, but all offers are such that local slopes are (approximately) the same. This is however very similar to the globally linear supply case. Alternatively, the piece-wise linear schedule is characterized by very heterogeneous local slopes, but the dynamics of the market is such that demand ends up intersecting always the same piece of the supply schedule - i.e. any given hourly market is virtually always cleared by the same plant. Heterogeneity across generating companies makes the latter case sound more likely. All of this suggests to view the proposed model as a piece-wise linear function. Variations of demand and the structure of supply are supposed to interact in such a way as to

¹²The different order system adopted by the APX may be the reason behind the comparatively poor performance of the $1/P$ volatility-price scaling for that market. This can be seen as a nice example of how much institutions matter in shaping the properties of market outcomes.

determine some high persistency as regards which plant is assigned the market-clearing role in a given hourly market.¹³

3.1.2 Economic meaning of the intercept and slope parameters

The proposed (piece-wise) linear model of the supply schedule is characterized by just two parameters (for each piece of the curve): intercept and slope. The former defines the so-called supply effect; the latter, multiplied by ΔQ_t , measures the demand effect. The evidence of $1/P$ scaling can be reproduced by assuming that such effects are additive.

Consider a piece-wise linear supply function. At any time, the relevant intercept parameter is the intercept of the linear function which the inelastic demand schedule intersects. This can fluctuate over time for a number of reasons.

A first mechanism is related to the evidence reported by Guerci et al. (2005). According to their analysis of the Spanish market, a certain fraction of power is supplied by the so-called *passive sellers* - i.e. by sellers posting an offered price equal to zero. Their orders can be seen as a sort of market orders. Those who post positive prices are instead called *active sellers*.

If either the number of passive sellers or the capacity offered by them (or both) change over time, parallel shifts of the supply function occur. Holding constant the offer prices and the capacities of active sellers, shifts due to passive agents modify the value of the intercept of each linear piece composing the market supply schedule, while leaving the slopes unaffected. Specifically, intercepts decrease following an increase in the total capacity supplied by passive sellers. The dynamics of entry, exit and capacity withholding by passive agents are probably driven by strategic considerations. An expected increase in demand would induce, on one hand, passive sellers to enact a supply reduction, adding to the demand effect and yielding higher profits for the benefit of all inframarginal sellers through a higher market-clearing price; on the other hand, outsider generating companies may decide to enter the market.

In markets with a considerable fraction of hydropower generation - such as the NordPool - parallel shifts of the supply schedule may be due to increases or decreases in the level of water reservoirs. These are slowly-varying forces, mostly dependent on meteorological and climatic factors, way beyond the control of power generating companies, and largely exogenous to demand fluctuations. The entailed supply effect has no strategic content.

Capacity withholding by *active agents*, too, may drive overall shifts in the supply schedule. However, shifts of this type can also be due to changes in the markups charged by generating companies on operating costs - in auction-theoretic terms, to a varying amount of *bid shading*. Both capacity withholding by active agents and the dynamics of bid shading are supposed to affect the local slopes of the supply function, too.

The evidence of $1/P$ volatility-price scaling is consistent with a supply schedule characterized by stochastic intercept α_t and constant slope β . Under such a supply model, most of the strategically-driven action behind the dynamics of wholesale electricity prices occurs by way of entry, exit and capacity changes by agents posting market orders. These phenomena are indeed captured by intercept shifts. Changes in the slope are instead assumed away: the demand effect is functionally stable over time. In light of the foregoing discussion, strategic moves by active agents, which would spoil the time constancy of β at the market-clearing volume level, play a relatively minor role. Moreover the time constancy of the slope requires that in any given hourly market, demand and supply always intersect in the same branch of the supply schedule. This can happen if the dynamics of passive capacity is positively correlated

¹³In such a case, the assumptions of continuity and differentiability of the supply function have to be weakened and given a local interpretation.

with the dynamics of demand, and at the same time if the structure of the supply schedule is stable in regard to the merit order of companies posting limit orders, their offered capacities, and their markups. For instance, an expected increase in demand could create incentives for entry of new passive sellers and for capacity withholding by extant ones. These effects imply, respectively, a right-ward and a left-ward shift of the supply schedule. Because demand would increase (i.e. the inelastic demand schedule would shift to the right), intersection between demand and supply curves keeps occurring in the same piece of the supply schedule if the right-ward shift of the supply curve overwhelms the left-ward shift - i.e. if a net increase of passive capacity occurs.

3.2 Distributions of price and demand dynamics

The empirical evidence reported in this paper shows that the kurtosis of filtered electricity price changes is higher than for filtered changes in electricity demand. This section investigates upon the properties of supply which, given a random and inelastic demand, yield the observed asymmetry between the distributions of the filtered ΔP_t and ΔQ_t . The forthcoming analysis assumes a supply schedule model characterized by the following features:

1. the inverse supply function is piece-wise linear with random intercepts and constant slopes;
2. shifts in the intercept term (supply effect) are correlated with demand dynamics;
3. intercept shifts are drawn from a zero-mean, symmetric distribution.

Linearity with random intercept and constant slope has been shown before to be sufficient for the $1/P$ scaling to emerge. The correlation between the supply effect and demand dynamics captures the aggregate impact of two ways through which suppliers can react to expected growth in demand: capacity withholding by existing generating companies, and entry of new capacity. The former determines an increase of the supply intercept, the latter a downward shift. Therefore we expect $\Delta\alpha_t$ to be negatively correlated with ΔQ_t whenever entry is quantitatively more relevant than capacity withholding. Otherwise, a positive correlation regime sets in.

The following research questions are asked. First, which properties of the intercept-shift process $\Delta\alpha_t$ and which values of β are consistent with the observed variance and kurtosis of ΔP_t and ΔQ_t ? Second, and more generally: under which conditions the properties of the intercept-shift variable $\Delta\alpha_t$ determine a divergence between the tail properties of price and demand dynamics? The forthcoming two subsections provide some answers.

3.2.1 Implied variance of the intercept process.

Let $\Delta\alpha_t \equiv \alpha_t - \alpha_{t-1}$ and ΔQ_t be random variables, which distributions are symmetric around a zero mean. Let v_x be the variance of a random variable x , and $\gamma_{x,y}$ the covariance between x and y . Under the assumed supply model, the variance of ΔP_t reads

$$v_{\Delta P} = v_{\Delta\alpha} + \beta^2 v_{\Delta Q} + 2\beta\gamma_{\Delta\alpha, \Delta Q} \quad (13)$$

The variance of volume dynamics, the variance of the intercept shift process, and the covariance between intercept shifts and volume growth - all have a positive impact on the

variance of ΔP . Holding $v_{\Delta P}$, $v_{\Delta Q}$ and $\gamma_{\Delta\alpha, \Delta Q}$ fixed, the local slope of the relationship between β and $v_{\Delta\alpha}$ reads

$$\frac{dv_{\Delta\alpha}}{d\beta} = -2(\beta v_{\Delta Q} + \gamma_{\Delta\alpha, \Delta Q}) \quad (14)$$

This can be seen as the slope of a level curve - the locus of all pairs $(v_{\Delta\alpha}, \beta)$ such that Eq. 13 holds. Such a slope is negative if $\gamma_{\Delta\alpha, \Delta Q} > -\beta v_{\Delta Q}$. This occurs, for instance, when capacity withholding reacts to demand growth more than entry. In such a case, steeper supply functions must be characterized by low volatility of their parallel shifts.

A steep slope could emerge in a market where very heterogeneous ask prices are submitted for very small capacity ranges. This given, if supply shifts are negatively correlated with demand dynamics, the proposed model implies a supply function which experiences very volatile shifts. This seems to be a good description of the dynamics of daily auctions. On the contrary, when any given offered price is posted for a wide capacity range, and offers are rather homogeneous across companies, β is small, and under a negative correlation regime one should observe rather stable, non-volatile parallel shifts in the supply schedule. This fits better for night auctions.

Notice that if $\beta = 0$, then $v_{\Delta\alpha} = v_{\Delta P}$; and that $v_{\Delta\alpha} = 0$ if

$$\beta = \bar{\beta} \equiv -\rho_{\Delta\alpha, \Delta Q} \sqrt{\frac{v_{\Delta\alpha}}{v_{\Delta Q}}} \pm \sqrt{\rho_{\Delta\alpha, \Delta Q}^2 \frac{v_{\Delta\alpha}}{v_{\Delta Q}} + \frac{v_{\Delta P}}{v_{\Delta Q}}}$$

where ρ denotes a correlation coefficient. Values of $\beta > \bar{\beta}$ would imply negative variance of $\Delta\alpha$, which is impossible. Hence, the range of feasible β values is $[0, \bar{\beta}]$.

3.2.2 Explaining kurtosis patterns.

Let k_x be the kurtosis of a random variable x . Under the assumed model, the kurtosis of the price change variable ΔP reads:

$$k_{\Delta P} = \frac{v_{\Delta\alpha}^2 k_{\Delta\alpha} + \beta^4 v_{\Delta Q}^2 k_{\Delta Q} + 6\beta^2 \gamma_{\Delta\alpha^2, \Delta Q^2} + 4\beta \gamma_{\Delta\alpha^3, \Delta Q} + 4\beta^3 \gamma_{\Delta\alpha, \Delta Q^3}}{v_{\Delta P}^2} \quad (15)$$

This complicated expression can be solved for the difference $k_{\Delta P} - k_{\Delta Q}$, which empirically is known to be positive:

$$k_{\Delta P} - k_{\Delta Q} = \frac{1}{v_{\Delta Q}^2} [a_1 \beta^{-1} + a_2 \beta^{-2} + a_3 \beta^{-3} + a_4 \beta^{-4}] \quad (16)$$

where:

$$a_1 \equiv 4(\gamma_{\Delta\alpha, \Delta Q^3} - k_{\Delta P} v_{\Delta Q} \gamma_{\Delta\alpha, \Delta Q}) \quad a_2 \equiv 2(3\gamma_{\Delta\alpha^2, \Delta Q^2} - 2k_{\Delta P} \gamma_{\Delta\alpha, \Delta Q}^2 - k_{\Delta P} v_{\Delta\alpha} v_{\Delta Q})$$

$$a_3 \equiv 4(\gamma_{\Delta\alpha^3, \Delta Q} - k_{\Delta P} v_{\Delta\alpha} \gamma_{\Delta\alpha, \Delta Q}) \quad a_4 \equiv v_{\Delta\alpha}^2 (k_{\Delta\alpha} - k_{\Delta P})$$

The gap between kurtosis values of price and volume dynamics depends negatively on the correlation between the supply effect and demand fluctuations:

$$\frac{\partial(k_{\Delta P} - k_{\Delta Q})}{\partial\gamma_{\Delta\alpha,\Delta Q}} = -\frac{4k_{\Delta P}v_{\Delta P}}{\beta^4v_{\Delta Q}^2} < 0 \quad (17)$$

A high kurtosis gap can thus be due to a strongly negative correlation between supply curve shifts and demand growth, as is the case when the aggregate effect of entry and exit by passive sellers is greater than the aggregate effect of changing incumbents' capacities. On the contrary, the impact of the kurtosis of intercept shifts is positive:

$$\frac{\partial(k_{\Delta P} - k_{\Delta Q})}{\partial k_{\Delta\alpha}} = \frac{v_{\Delta\alpha}^2}{\beta^4v_{\Delta Q}^2} > 0 \quad (18)$$

This suggests that a large kurtosis gap may occur when the probability of very large shifts in the supply function is non-negligible. The relationship with the variance of demand growth is positive only if $k_{\Delta\alpha} > \frac{v_{\Delta P}}{v_{\Delta\alpha}}k_{\Delta P}$ - i.e. if intercept fluctuations are fat-tailed enough:

$$\frac{\partial(k_{\Delta P} - k_{\Delta Q})}{\partial v_{\Delta\alpha}} = -\frac{2(k_{\Delta P}v_{\Delta P} - k_{\Delta\alpha}v_{\Delta\alpha})}{\beta^4v_{\Delta Q}^2} \quad (19)$$

The gap between kurtosis of price and volume fluctuations is increasing in the value of the higher order covariances, as it can be easily verified. Finally, notice that the kurtosis gap vanishes as β grows large:

$$\lim_{\beta \rightarrow \infty} k_{\Delta P} - k_{\Delta Q} = 0 \quad (20)$$

However, the relationship between $k_{\Delta P} - k_{\Delta Q}$ and β - a quartic function - is not in general monotonic.

To sum up, the kurtosis gap can be reproduced by assuming: (i) the impact of entry and exit, at the aggregate market level, is higher than the impact of capacity withholding (negative correlation between $\Delta\alpha$ and ΔQ); (ii) volatile demand growth and volatile shifts in the supply schedule are associated (positive high order correlations, such as $\gamma_{\Delta\alpha^2,\Delta Q^2}$); and (iii) some coordination of individual strategic moves exists, such that the supply schedule experiences either very small or very large parallel shifts (large kurtosis of $\Delta\alpha$).

Case: linear supply with non-stochastic parameters. In order to fully appreciate the role of a stochastic supply effect, let us study the special case of a linear model with constant parameters; namely

$$P_t = \alpha + \beta Q_t \quad (21)$$

In this case, $v_\alpha = 0$, implying $v_{\Delta P} = \beta^2v_{\Delta Q}$, and $k_{\Delta P} = k_{\Delta Q}$. Because ΔQ_t is symmetric by assumption, also odd moments of price and demand changes are equal. As a conclusion, under a linear supply function with non-stochastic parameters, the distributions of ΔP_t and ΔQ_t are identical up to a constant factor β .¹⁴

As another way to see this, recall that the equality

$$\Delta P_t = \beta \Delta Q_t \quad (22)$$

¹⁴It is worth noting that such a conclusion is based on the comparison between kurtosis indicators, which are a-dimensional, and between variances of ΔP_t and of $\beta \Delta Q_t$, which are expressed in the same units by a suitable choice of β .

implies, for $\Delta Q \sim f_{\Delta Q}$, that

$$f_{\Delta P}(x) = \frac{1}{\beta} f_{\Delta Q}(x) \quad (23)$$

$\forall x$ in the relevant domain. The observed asymmetry between tail behaviors of the distributions of price and demand dynamics cannot be reproduced, unless α and/or β are random variables. Introducing randomness in α is therefore a way to reproduce the different tail behavior of changes in prices and demand.

3.3 Autocorrelation patterns

The previous empirical analysis has shown that the weekly pattern of electricity prices, measured by its autocorrelation at 7 days, is stronger than for demand dynamics. Let $\rho_x(k)$ be the autocorrelation coefficient of a random variable x at lag k , and $\rho_{x,y}(k)$ the cross-correlation between x and y at lag k . Furtherly assume that all variables of interest are covariance stationary. Under the assumed supply model, the following holds:

$$\rho_{\Delta P}(k) = v_{\Delta P}^{-1} \left[v_{\Delta\alpha} \rho_{\Delta\alpha}(k) + \beta^2 v_{\Delta Q} \rho_{\Delta Q}(k) + 2\beta \sqrt{v_{\Delta\alpha} v_{\Delta Q}} \rho_{\Delta\alpha, \Delta Q}(k) \right] \quad (24)$$

The observed regularity $\rho_{\Delta P}(7) > \rho_{\Delta Q}(7)$ requires that the following condition is satisfied:

$$(\beta^2 v_{\Delta Q} - v_{\Delta P}) \rho_{\Delta Q}(k) + v_{\Delta\alpha} \rho_{\Delta\alpha}(k) + 2\beta \sqrt{v_{\Delta\alpha} v_{\Delta Q}} \rho_{\Delta\alpha, \Delta Q}(k) > 0$$

Plugging Eq. 13 into the above and doing some algebra yields:

$$v_{\Delta\alpha} (\rho_{\Delta\alpha}(7) - \rho_{\Delta Q}(7)) + 2\beta \sqrt{v_{\Delta\alpha} v_{\Delta Q}} (\rho_{\Delta\alpha, \Delta Q}(7) - \rho_{\Delta\alpha, \Delta Q}(0)) > 0 \quad (25)$$

This condition indicates that the autocorrelation structure of intercept shifts ($\rho_{\Delta\alpha}$) and its cross-correlation with changes in demand ($\rho_{\Delta\alpha, \Delta Q}$) play an important role in driving the weekly pattern observed in wholesale electricity price dynamics. The former indicates persistency of the supply effect. The latter denotes the existence of a strategic linkage between the supply effect (e.g. capacity withholding, entry, exit) and demand dynamics.

The observed autocorrelations pattern depends on (i) whether $\Delta\alpha$ is more persistent than ΔQ (first term of the summation in Eq. 27); and on (ii) how slowly the correlation between $\Delta\alpha$ and ΔQ decays over time (2nd term). The observed evidence is more likely to be verified if demand growth and the supply effect are very volatile, and when β is high.

The interaction between supply shifts and the expected demand dynamics is the truly interesting economic aspect of the dynamics of wholesale electricity auctions. Were suppliers not to react to either expected or past realized demand dynamics, the difference between magnitudes of price and demand growth autocorrelations would be milder. For instance, were $\Delta\alpha$ deterministic and not correlated with ΔQ , Eq. 27 would be identically equal to 0, i.e. no gap between autocorrelations would be observed.

To conclude, the existence of an autocorrelation gap indicates that the strategic response to expected demand dynamics is persistent and periodic (high autocorrelation of $\Delta\alpha$ at weekly lags); and (ii) such strategic moves and the expectations they are based upon are linked to the past dynamics of demand, in an adaptive fashion (high lagged correlation between $\Delta\alpha$ and ΔQ).

4 Concluding remarks

In this paper, the statistical properties of day-ahead electricity prices have been studied through a new approach. The main idea is that, given an inelastic demand schedule and a uniform-price auction format, one can derive the properties of price growth through the statistical regularities observed in power demand growth, and that such properties vary under different assumptions about the supply curve. Intuitively, only some supply curve models are consistent with the existing joint evidence on prices and volumes.

Consistently, the properties of day-ahead electricity price and volume growth rates have been jointly analyzed. More in detail, it has been shown that serial correlations of the growth of prices are stronger than for volume growth; that price growth is more heavy-tailed than volume growth; and that the conditional standard deviation of price growth decays like the reciprocal of the price level ($1/P$ scaling). The latter fact can be explained by a piece-wise linear supply function, with stochastic local intercepts and constant local slopes. The other two facts have been interpreted in light of the proposed model. The analysis shows that most of the action behind day-ahead electricity price dynamics is due to entry, exit, and capacity withholding by the so-called passive sellers, whereas changes in the merit order of active sellers plays a minor role. In models in which generating companies respond to changing demand conditions in heterogeneous ways (entry, exit, capacity withholding), the kurtosis gap can be reproduced by assuming (i) that the aggregate impact of entry and exit is quantitatively greater than the effect of capacity withholding; (ii) that expectations of a volatile demand growth are associated with volatile shifts in the supply schedule; and (iii) that individual reactions are coordinated in such a way, that parallel shifts in the supply schedule are either very small, or very large. The observed gap between serial correlations of price growth and demand growth is due to persistency in the strategic responses to expected demand growth, and to adaptive behavior in the process of expectation formation.

Studying the *interdependence* between demand and supply is therefore a key to a deeper understanding of how day-ahead uniform-price auctions for electricity work. Availability of data about individual offers - i.e. about the empirical market curves - will further enhance such understanding, and will allow to validate the proposed model and to amend it for consistency with real-world properties.

As a long-term objective of this research, one would like to build an auction model capable to predict the patterns of entry, exit and capacity variations which seem to be driving the observed empirical regularities. The empirically-based model put forward in this paper can be seen as an aggregate perspective to understanding the electricity market dynamics. Extensions of models such as von der Fehr and Harbord (1993) and Crampes and Crete (2003) may provide the necessary microfoundations and shed light on the empirical strengths and weaknesses of the existing, axiomatic auction-theoretic models of electricity pools.

References

- [1] F.L. Alvarado, R. Rajaraman (2000). Understanding Price Volatility in Electricity Markets. *Proceedings of the 33rd Hawaii International Conference on System Sciences*.
- [2] F. Bellini (2002). Empirical Analysis of Electricity Spot Prices in European Deregulated Markets. *Quaderni Ref.* 7/2002.

- [3] S. Blumsack, D. Perekhodtsev, L.B. Lave (2002). Market Power in Deregulated Wholesale Electricity Markets: Issues in Measurement and the Cost of Mitigation. *Electricity Journal*, November: 11-24.
- [4] G. Bottazzi (2004), *Subbotools: A Reference Manual*. LEM Working Paper 2004/14, S.Anna School of Advanced Studies, Pisa.
- [5] G. Bottazzi, S. Sapio, A. Secchi (2005). Some Statistical Investigations on the Nature and Dynamics of Electricity Prices. *Physica A* 355(1): 54-61.
- [6] G. Bottazzi, A. Secchi (2003). Why are Distributions of Firm Growth Rates Tent-Shaped?, *Economics Letters* 80: 415–420.
- [7] G. Bottazzi, A. Secchi (2005). Explaining the Distribution of Firms Growth Rates. *RAND Journal of Economics*, forthcoming.
- [8] H. Bystroem (2001). Extreme Value Theory and Extremely Large Electricity Price Changes. Mimeo.
- [9] L. Cavallo, S. Sapio, V. Termini (2005). Market Design and Electricity Spot Prices: Empirical Evidence from NordPool and California Price Crises. Mimeo.
- [10] S.-J. Deng, W. Jiang, Z. Xia (2002). Alternative Statistical Specifications of Commodity Price Distribution with Fat Tails. *Advanced Modeling and Optimization* 4: 1–8.
- [11] F.X. Diebold, L.E. Ohanian, J. Berkovitz (1997). Dynamic Equilibrium Economies: A Framework for Comparing Models and Data. Federal Reserve Bank of Philadelphia, working paper No. 97-7.
- [12] E. Eberlein, G. Stahl (2003). Both Sides of the Fence: a Statistical and Regulatory View of Electricity Risk. *Energy and Power Risk Management* 8: 34–38.
- [13] H. Geman, A. Roncoroni (2002). Understanding the Fine Structure of Electricity Prices. Forthcoming, *Journal of Business*.
- [14] J.-M. Glachant, D. Finon, eds. (2003). *Competition in European Electricity Markets: A Cross Country Comparison*. Cheltenham: Edward Elgar.
- [15] R. Green (2002). Market Power Mitigation in the UK Power Market. Paper prepared for the conference Towards a European Market for Electricity, organized by the SSPA, Italian Advanced School of Public Administration, Rome, June 24-5, 2002.
- [16] E. Guerci, S. Ivaldi, S. Pastore, S. Cincotti (2005). Modeling and Implementation of an Artificial Electricity Market Using Agent-Based Technology. *Physica A* 355: 69-76.
- [17] E. Guerci, S. Ivaldi, P. Magioncalda, M. Raberto, S. Cincotti (2006). Modeling and Forecasting of Electricity Prices Using an ARMA-GARCH Approach. Forthcoming, *Quantitative Finance*.
- [18] G.L.F. Holburn, P.T. Spiller (2002). Institutional or Structural: Lessons from International Electricity Sector Reforms. In E. Brousseau, J.M. Glachant (eds.), *The Economics of Contracts: Theories and Applications*. Cambridge, MA: Cambridge University Press.

- [19] P.L. Joskow (1996). Introducing Competition into Regulated Network Industries: from Hierarchies to Markets in Electricity. *Industrial and Corporate Change* 5: 341–382.
- [20] C.R. Knittel, M. Roberts (2005). An Empirical Examination of Deregulated Electricity Prices. POWER WP-087, University of California Energy Institute. *Energy Economics* 27(5): 791-817.
- [21] T.J. Kozubowski, K. Podgorski (2001). Asymmetric Laplace Laws and Modeling Financial Data. *Mathematical and Computational Modelling* 34: 1003-1021.
- [22] Y. Lee, L.A.N. Amaral, D. Canning, M. Meyer, H.E. Stanley (1998). Universal Features in the Growth Dynamics of Complex Organizations. *Physical Review Letters* 81, 3275-3278.
- [23] F.A. Longstaff, A. W. Wang (2002). Electricity Forward Prices: a High-Frequency Empirical Analysis. UCLA Anderson School working paper.
- [24] J.J. Lucia, E. Schwartz (2002). Electricity Prices and Power Derivatives: Evidence from the Nordic Power Exchange. *Review of Derivative Research* 5(1): 5-50.
- [25] J.L. McCauley (2004). *Dynamics of Markets. Econophysics and Finance*. Cambridge: Cambridge University Press.
- [26] L. Meeus (2005). Pricing Electric Energy Auctions with Block Orders. Presented at the ETE Workshop on Market Modelling of the Central Western European Electricity Market - K.U. Leuven, Belgium, September 2005.
- [27] T. Mount (2000). Strategic Behavior in Spot Markets for Electricity when Load is Stochastic *Proceedings of the 33rd International Hawaii Conference on System Sciences*.
- [28] D. Newbery (2002). Issues and Options for Restructuring Electricity Supply Industries. Cambridge Working Papers in Economics 0210, Department of Applied Economics, University of Cambridge.
- [29] J. Nowicka-Zagrajek, R. Weron (2002). Modeling electricity loads in California: ARMA models with hyperbolic noise. *Signal Processing* 82(12): 1903-1915.
- [30] S. Sapio (2004). Market Design, Bidding Rules, and Long Memory in Electricity Prices. *Revue d'Economie Industrielle* 107: 151-170.
- [31] S. Sapio (2005). *The Nature and Dynamics of Electricity Markets*. Doctoral dissertation, Sant'Anna School of Advanced Studies, Pisa.
- [32] I. Simonsen (2005). Volatility of Power Markets. *Physica A* 355: 10-20.
- [33] M.H.R. Stanley, L.A.N. Amaral, S.V. Buldyrev, S. Havlin, H. Leschorn, P. Maass, M.A. Salinger, H.E. Stanley (1996). Scaling Behavior in the Growth of Companies, *Nature* 379: 804.
- [34] M.F. Subbotin (1923). On the Law of Frequency of Errors. *Matematicheskii Sbornik* 31: 296–301.
- [35] N. von der Fehr, D. Harbord (1993). Spot Market Competition in the UK Electricity Industry. *Economic Journal* 103: 531–546.

- [36] R. Weron (2002). Measuring Long-Range Dependence in Electricity Prices. In H. Takayasu (ed.), *Empirical Science of Financial Fluctuations*. Springer, Tokyo: 110-119.
- [37] R. Weron, M. Bierbrauer, S. Truck (2004). Modeling Electricity Prices: Jump Diffusion and Regime Switching. *Physica A* 336: 39–48.
- [38] M. West (1987). On Scale Mixtures of Normal Distributions. *Biometrika* 74: 646-648.
- [39] R. Wilson (2000). Architecture of Power Markets. *Econometrica* 70(4): 1299-1344.
- [40] F.A. Wolak, R.H. Patrick (1997). The Impact of Market Rules and Market Structure on the Price Determination Process in the England and Wales Electricity Market. Working Paper PWP-047, University of California Energy Institute.
- [41] A. Worthington, A. Kay-Spradley, H. Higgs (2002). Transmission of Prices and Price Volatility in Australian Electricity Spot Markets: a Multivariate GARCH Analysis. Queensland University of Technology, School of Economics and Finance Discussion Paper No. 114.