Innovation processes in logically constrained time series

Christoph Möller, Svetlozar T. Rachev, Young S. Kim and Frank J. Fabozzi

Abstract Capturing the relevant aspects of phenomena in an econometric model is a fine art. When it comes to the innovation process a trade of between a suitable process and its mathematical implications has to be found.
In many phenomena the likelihood of extreme events plays a crucial role. At the same time, classical extreme value theory is based on assumptions that cannot logically be drawn for the phenomenon in question. In this paper, we exemplify the fitness of tempered stable laws to capture both the probability of extreme events, and the relevant boundary conditions in a back-coupled system, the German balancing energy demand.

Key words: tempered stable, SARIMA, innovation process, balancing energy

Christoph Möller
Department of Statistics, Econometrics and Mathematical Finance, School of Economics and Business Engineering, University of Karlsruhe and KIT, Geb. 20.12, 76131 Karlsruhe, Germany e-mail: moeller@statistik.uni-karlsruhe.de

Svetlozar T. Rachev
Chair-Professor, Chair of Statistics, Econometrics and Mathematical Finance, School of Economics and Business Engineering, University of Karlsruhe and KIT, and Department of Statistics and Applied Probability, University of California, Santa Barbara, and Chief-Scientist, FinAnalytica INC, Geb. 20.12, 76131 Karlsruhe, Germany e-mail: rachev@statistik.uni-karlsruhe.de

Young Shin Kim
Department of Statistics, Econometrics and Mathematical Finance, School of Economics and Business Engineering, University of Karlsruhe and KIT, Geb. 20.12, 76131 Karlsruhe, Germany e-mail: kim@statistik.uni-karlsruhe.de

Frank J. Fabozzi
Professor in the Practice of Finance, Yale School of Management New Haven, CT USA e-mail: frank.fabozzi@yale.edu
[*]Svetlozar T. Rachev gratefully acknowledges research support by grants from Division of Mathematical, Life and Physical Sciences, College of Letters and Science, University of California, Santa Barbara, the DFG and the DAAD. The authors thank Prof. Gennady Samorodnitsky for his help in formulating the problem of TID distributions and for his fruitful comments and suggestions.

1 Introduction

Classical time series analysis assumes a normal distribution. This assumption is justified by the central limit theorem (CLT) which states that the infinite sum of random numbers drawn from any distribution with finite first and second moments converges to the normal distribution. The global financial crisis that began in late 2007, however, has demonstrated how problematic the assumption of normal innovations or a Gaussian copular is when it comes to evaluating risk. There is a rich literature proposing alternative modeling approaches (see, for example, Rachev and Mittnik (2000) and the references within).

In this paper we will analyze a non-financial seasonal autoregressive integrated moving average (SARIMA) time series whose innovation process is clearly in the domain of the CLT due to physical boundary conditions. Nonetheless, we will demonstrate normal innovations to be inadequate in this setting due to the presence of heavy-tailed distributions in the data. Furthermore, we will address the issue of robustness of the Kolmogorow-Smirnow test and the Anderson-Darling test when testing heavy-tailed distributions.

The remainder of this paper is organized as follows. Section 2 reviews classical linear time series analysis: the non-seasonal autoregressive integrated moving average (ARIMA) and SARIMA models. Different distribution classes for the corresponding innovation process are discussed in Section 3 along with distribution tests. Section 4 introduces the balancing energy demand time series and the relevant boundary conditions. In Section 5 we describe the SARIMA model used to filter the innovations that are subsequently analyzed for their distribution in Section 6. Section 7 concludes and summarizes our results.

2 Time series models

In econometrics, a linear time series regression can be regarded as a fundamental building block. We restrict our overview here to the case of discrete time series models. One widely known example is the autoregressive (AR) model capturing serial dependence in an analyzed time series. In this model, the realized time series x_t of a random variable X_t is generated from a noise process ε_t, in its standard form the Gaussian white noise process, and a regression on the past realizations of the process. One example would be Brownian motion w_t in discrete time; however, this process is not stationary and may grow without limits. Nevertheless, the process is transformed into a stationary process $(w_t - w_{t-1})$ by differencing, which in turn is expanded into the original Brownian motion by summation or integration. So Brownian motion is a simple example of an integrated stationary AR-process.

If one replaces each past realization in the AR-process by the lagged process, it is transformed into a possibly infinite regression on past realizations of the noise process, the so-called moving average (MA) process. Obviously this transformation is reversible and so any such MA-process can be transformed into a possibly infinite

AR-process. Additionally, Wold's decomposition theorem states that any stationary time series may be represented by a possibly infinite MA-process, or equivalently AR-process. This explains why the combination of the autoregressive moving average (ARMA)-process has become the standard approach in time series modeling. The practical advantage over the strict AR- and MA-processes is a parsimonious usage of parameters. This approach is expanded to an even wider range of non-stationary time series the ARIMA model, when including the appropriate degree of differencing to obtain a stationary time series.
Again the consideration of parameter parsimony leads to a further specification, the SARIMA model. This seasonal ARIMA model is a serial connection of ARIMA models. Here, one ARIMA model captures a serial dependence structure over cyclical periods (e.g., one year), and a further ARIMA model captures the dependence on a sub-cycle horizon (Box and Jenkins (1970)). Using the lag-operator L and the difference-operator ∇, the model can be expressed as in Equation (1):

$$\phi_p(L)\Phi_P(L^s)\nabla^d\nabla^D x_t = \theta_q(L)\Theta_Q(L^s)\varepsilon_t \tag{1}$$

Here s is the seasonality, p and P are the orders of the AR-polynomials ϕ, Φ and q and Q are the orders of MA-polynomials θ and Θ, and d and D are the appropriate numbers of differencing to obtain stationarity.
Therefore the model is also abbreviated as $SARIMA(p,d,q)\times(P,D,Q)_s$. The crucial difference to the regular $ARIMA(p,d,q)$ model can be best seen by expanding the multiplicative model into its additive form as in Equation (2):

$$\begin{aligned}(1-\phi_1 L-\phi_2 L^2-\ldots-(\phi_s+\Phi_1)L^s-(\phi_{s+1}-\phi_s\Phi_1)L^{s+1}-\ldots)\nabla^d\nabla^D x_t\\ =(1+\theta_1 L+\theta_2 L^2+\ldots+(\theta_s+\Theta_1)L^s+(\theta_{s+1}+\theta_s\Theta_1)L^{s+1}+\ldots)\varepsilon_t\end{aligned} \tag{2}$$

So as an example, a $SARIMA(1,0,0)\times(1,0,0)_{12}$ is equivalent to an $ARIMA(13,0,0)$ with the constraint $\phi_{13}=-\phi_1\phi_{12}$.
We conclude this section with a brief overview of model identification techniques that were established by Box and Jenkins (1970). The first condition a time series needs to fulfill in order to adapt a linear model is stationarity. This condition is fulfilled if all roots of the polynomial $\phi(L)$ lie outside the unit circle. Consequently, the hypothesis of a root of value one is tested (i.e., the unit root test). Should this hypothesis be rejected, the time series may be regarded as stationary, otherwise a higher order of differencing has to be considered. The work of Dickey and Fuller (1979) lead to the augmented Dickey-Fuller test (ADF-test) as an unbiased unit root test.
Once the appropriate degree d of differencing is selected, inspection of the sample autocorrelation function (SACF) and partial autocorrelation function (SPACF) will help to select the order of the AR and MA polynomials. The SACF will cut off at the lag q of a $MA(q)$-process, while the SPACF will cut off at the lag p of an $AR(p)$-process. Unfortunately, the SACF will only gradually tail off for an AR-process as

will the SPACF for a MA-process. Consequently, an *ARMA*-process will lead to a tailing off of both SACF and SPACF. Additionally, because it is calculated on a limited sample, this method is also prone to statistical fluctuations. Nonetheless, sudden drops in SACF and SPACF give an indication of appropriate choices for the parameters p and q of an *ARMA*-process.
Finally model selection criteria such as the Akaike information criterion (AIC) or Bayesian information criterion (BIC) should be used as a guide to avoid the pitfalls of overparameterized modeling.

3 Innovations

The main purpose of modeling time series is to identify and understand underlying characteristics and then be able to apply this understanding to predict the future. While a specification of the appropriate parameters of a SARIMA$(p,d,q) \times (P,D,Q)_s$ captures information on the location parameter of a future realization, the identification of the distribution of the innovation process ε_t is crucial in understanding the risk that goes along with anticipating this realization.
Among all distributions, the class of α-stable distributions stands out as the distributional class any appropriately scaled sum of independent and identically distributed (i.i.d.) random variables will converge to. In the context of time series modeling, this property is of particular interest as the innovations in each time step are considered to be the sum of many independent events from the same set of plausible events (i.e., the sum of i.i.d. random variables).
The normal distribution is the most well-known member of the class of α-stable distributions, and for this distribution the convergence property is known as the CLT. Relaxing the assumption of existing first and second moments, the CLT is generalized. The limiting distribution is now not generally the normal distribution, but, instead, a member of the class of α-stable distributions. This is known as the general central limit theorem (GCLT). In the following we will refer to α-stable distributions as non-Gaussian to distinguish the class from its most well-known member. In general, for α-stable distributions there exists no probability density function in closed form and therefore it is expressed by its characteristic function as given by Equation (3):

$$\phi_{\text{stable}}(u;\alpha,\sigma,\beta,\mu) = E[e^{iuX}] \tag{3}$$
$$= \begin{cases} \exp\left(i\mu u - |\sigma u|^{\alpha}\left(1 - i\beta(\text{sign u})\tan\frac{\pi\alpha}{2}\right)\right), & \alpha \neq 1 \\ \exp\left(i\mu u - \sigma|u|\left(1 + i\beta\frac{2}{\pi}(\text{sign u})\ln|\text{u}|\right)\right), & \alpha = 1, \end{cases}$$

where

$$\text{sign } t = \begin{cases} 1, & t \geq 0 \\ 0, & t = 0 \\ -1, & t \leq 0 \end{cases}$$

In this parameterization, the four parameters $(\alpha, \beta, \sigma, \mu)$ have the following domain and interpretation:

- α: the index of stability or the shape parameter, $\alpha \in (0,2)$;
- β: the skewness parameter,$\beta \in [-1,+1]$;
- σ: the scale parameter, $\sigma \in (0,+\infty)$;
- μ: the location parameter, $\mu \in (-\infty,+\infty)$.

The α-stable distribution is a very rich class of distributions that can capture even extreme asymmetric and heavy-tailed structural conditions. However, in some problems the implication of infinite variance, or in the case of $\alpha \leq 1$ infinite mean and variance, is too strong an assumption. At the same time the empirical distribution might clearly show excess kurtosis or skewness. For these problems the class of tempered stable distributions has been proposed. (See Menn et al (2005), Kim et al (2008), and Menn and Rachev (2008)). In an adaption of the α-stable distribution, the tempered stable distribution has the characteristic function given by Equation (4) with one extra parameter.

$$\begin{aligned} \phi_{CTS}(u;\alpha,C,\lambda_+,\lambda_-,m) &= \exp(ium - iuC\Gamma(1-\alpha)(\lambda_+^{\alpha-1} - \lambda_-^{\alpha-1}) \\ &\quad + C\Gamma(-\alpha)((\lambda_+ - iu)^\alpha - \lambda_+^\alpha + (\lambda_- + iu)^\alpha - \lambda_-^\alpha)). \end{aligned} \tag{4}$$

Here m and C determine the location and scale as do μ and σ in the α-stable distribution. However, the skewness is parameterized by λ_+ and λ_-, at the same time allowing for a faster than α-stable decay in the tails. So a tempered stable distribution allows for the same flexibility at the center of the distribution as the α-stable distributions, combined with finite first and second moments. Therefore, the tempered stable distribution is in the domain of the CLT, though convergence to the Gaussian limit may be very slow due to the almost α-stable distribution in the center. So tempered stable distributions are a natural choice whenever the problem suggests a distribution that is prone to extreme events, but at the same time other considerations forbid infinite moments; that is, whenever the rate of convergence to the limiting Gaussian distribution at the infinite sum is too slow to justify a Gaussian model at the finite time horizon of interest.
The classical tempered stable (CTS) distribution has been introduced under different names in the literature, including the truncated Levy flight by Koponen (1995), the KoBoL distribution by Boyarchenko and Levendorskiĭ (2000), and the CGMY distribution by Carr et al (2002).
Having introduced the tempered stable distribution, we conclude this section with a review of methods to justify the introduction of additional parameters into a model by specifying a non-Gaussian innovation process. One widely used graphical method is the inspection of the quantile-quantile (QQ)-plot. Figure 1 shows an example of the quantiles of the fitted Gaussian-distribution versus the empiri-

cal quantiles of an observed innovations time series. Ideally, the QQ-plot would be a straight line of unit slope, and any deviation from this line is an indication of the chosen distribution not adequately describing the innovation time series. One should keep in mind that because it is based on a finite sample, the empirical quantiles will never stretch out to arbitrarily small or large quantiles.

Another less subjective method is testing the proposed distribution with a goodness-of-fit test. All the tests calculate a test statistic (see Table 1) to test the hypothesis of the innovations time series to be a realization of the fitted distribution. These statistics can then be compared to critical values and p-values can be calculated.

The Kolmogorov-Smirnov (KS) test is known for its critical values to be indepen-

Fig. 1 QQ-plot of heavy-tailed data under the Gaussian hypthesis

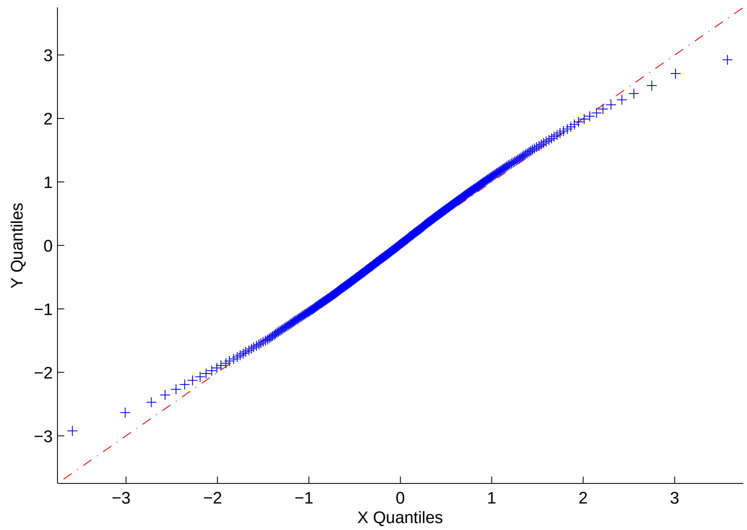

Table 1 Goodness-off-fit test and their statistic

test	statistic
KS	$\sqrt{n}\sup\lvert F_{emp.}(x) - F_{th.}(x)\rvert$
AD	$\sqrt{n}\sup\lvert\frac{F_{emp.}(x)-F_{th.}(x)}{\sqrt{F_{th.}(x)(1-F_{th.}(x))}}\rvert$
CvM	$n\int_{-\infty}^{\infty}(F_{emp.}(x) - F_{th.}(x))^2 dF_{th.}(x)$
AD^2	$n\int_{-\infty}^{\infty}(\frac{F_{emp.}(x)-F_{th.}(x)}{F_{th.}(x)(1-F_{th.}(x))})^2 dF_{th.}(x)$

dent of the tested distribution. Like the KS-test, the Anderson-Darling (AD) test is based on the supremum of the difference of the theoretical and empirical CDF, however, a weight assigned to each point assigns more weight to the tails of the distribution. In the case of the α-stable distribution, this feature is obviously of particular importance.

Both the KS-test and AD-test have to be used with caution when testing α-stable or tempered stable distributions. We illustrate this drawback in the following example. Suppose the underlying distribution is the CTS-$(0.3, 1, 0.5, 0.7, 0)$ shown in Figure 2. There is an almost vertical step in the CDF at the location parameter *m*. Because of the dominance of extreme events, the mean of a sample drawn from a heavy-tailed distribution will only converge slowly to the true mean of zero. Consequently, the steep increase is very likely to be dislocated in the empirical CDF. In combination with the almost vertical increase at the location parameter, this may result in large deviances between the empirical and theoretical CDF and a rejection of the distribution in the KS-test. This argument might even hold true for the AD-test where the weights are minimal at the location parameter. This behavior is particularly undesirable when analyzing innovation distributions where the true mean of zero is known and the location parameter is of minor importance. We therefore suggest a flexibility in adjusting the sample mean before a KS-test or an AD-test.

In contrast to the KS-test and AD-test, the Cramer-von Mises (CvM) test and the

Fig. 2 CDF of CTS-$(0.3, 1, 0.5, 0.7, 0)$

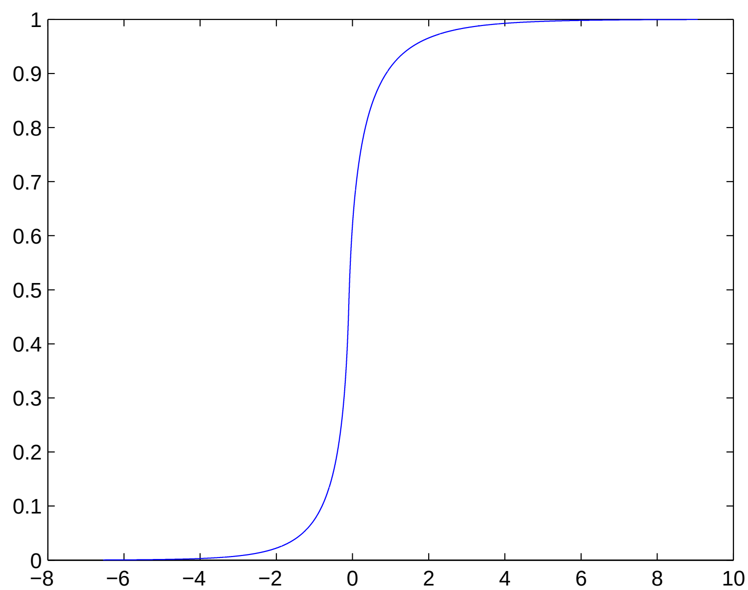

AD^2-test are based on the area between the theoretical and empirical CDF (see Table 1). The main advantage of these statistics is that they incorporate information about the total sample, and, in particular, are insensitive to a slight dislocation of the empirical CDF. The AD^2-test introduces a weighting scheme that focuses the test on the tails of the distribution. One major disadvantage of the CvM-test, AD^2-test and AD-test is that critical values are dependent on the analyzed distribution. One feasible method of obtaining critical values in these cases is to use Monte Carlo-based simulation based on the estimated parameter set (see Chernobai et al (2007)). Therefore, the tests are not available in common software, and in turn results are less transparent.

4 Market setting and data

Electricity is a special commodity as it is practically not storable. Consequently, in electricity grids supply and demand sides have to be in exact equilibrium at all times, otherwise a blackout will result. This task of active balancing is performed by the transmission system operator (TSO) who activates bids of up- or down-regulating energy to maintain equilibrium. In this context, it is important to distinguish between regulatory energy and balancing energy. Regulatory energy is contracted prior to the actual balancing action by the TSO to allocate the resources to be able to balance the grid. Balancing energy is the energy the TSO will settle with balancing responsible entities (BR) who caused a disturbance after the balancing action. The TSO performs this correction by calling up appropriate regulatory energy bids contracted prior to the actual balancing.

It is technically not feasible to balance disturbances of arbitrary sizes as the blackouts in northeast North America in 2003 or the blackout in central Europe in 2006 demonstrate. So balancing energy demand is constrained around zero; however, no exact boundary can be specified. In general, balancing energy has to compensate all unpredictable supply and demand shocks such as a power plant outage or a football match going into overtime. Such shocks are unknown in advance and not predictable over a time horizon of a day or even more. Within the German electricity market design, balancing energy may however also be used as a substitute for other energy contracts as long as this position is neither obvious nor excessive (see Bundesnetzagentur (2006)). This incentive leads to pronounced patterns and serial dependence (see Nailis and Ritzau (2006) and Möller et al (2009)).

In this paper we concentrate on the governing innovation distribution of balancing energy demand in Germany. To uncover this distribution, we therefore have to filter from the data predictable components. The analytical model used here is described in Möller et al (2009). For the German market two factors are identified. These factors are the gradient of grid load and an arbitrage incentive. Next we restate some of the results.

The gradient effect is modeled by the out of sample average, which in turn can be modeled by Equation (5) using the load ($L(t)$) as a factor.

$$q(\nabla L(t)) = q \cdot (\overline{L}_q(t) - \overline{L}_h(t)), \quad q = 0.42 \tag{5}$$

This model describes the quarter hourly pattern with a R^2 of 0.87. The effect is fully attributable to the discrepancy between the step function of load changes imposed by the traded contracts on the day-ahead market, and a rather smoothly changing load in reality. Such effects cannot be observed in most of the other European markets. The settlement periods of the day-ahead market and the balancing energy market are one and the same in these markets (see EU (2005)).

Figure 3 shows the estimated gradients joined to form one daily figure, wherein

Fig. 3 Comparison of model (light) and data (dark)using 2004 data.

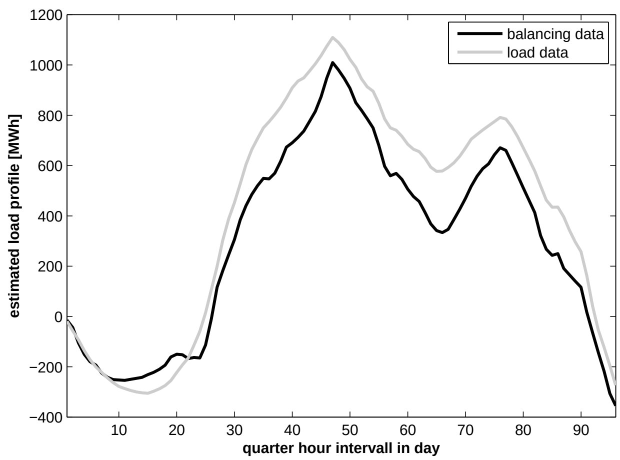

the missing information on the gradient between adjacent hours is estimated by the local average. This figure shows a remarkable resemblance of the typical German load profile, and illustrates the fitness of the model.

As for the second factor, there exist arbitrage incentives ($I(t), I_{tec}(t)$) to substitute electricity trades on the day-ahead market with strategic positions in the balancing energy market or avoid technical difficulties in plant operation. This incentive is driven by the day-ahead prices exceedence of a common price level and can be modeled as in Equation (6).

$$h(I(t), I_{tec}(t)) = a \cdot \left(\frac{2}{1 + b \cdot e^{-c \cdot I(t)}} - 1 \right) + I_{tec}(t) \quad \forall a \in \mathbb{R} b, c \in \mathbb{R}_+ \tag{6}$$

Figure 4 shows the fitted model together with the in-sample expectation values con-

Fig. 4 Factor model prediction and data

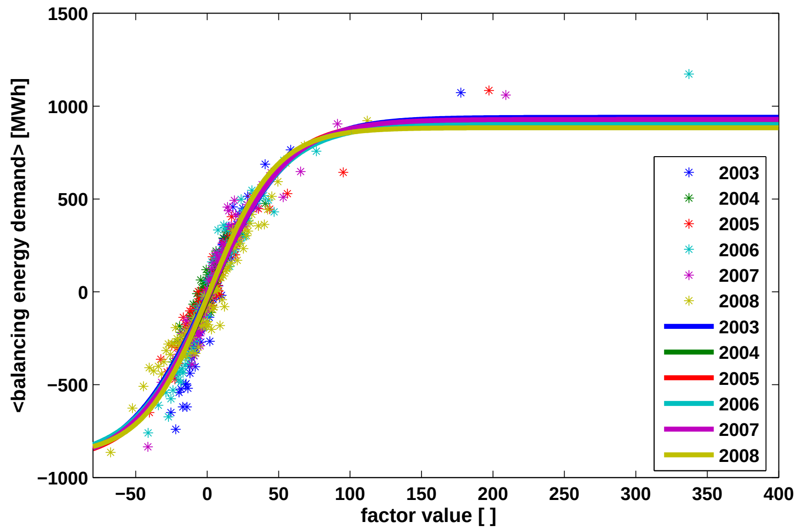

ditional on the factor values in different years. The data points represent conditional in-sample expectation values, while the fitted models are fitted to out-of-sample data, yielding the parameter estimates and R^2-values in Table 2. The model captures the dependence and saturation effects, and is compatible with a functional relation that is constant over time. The sample variance is reduced by 12% and 19% in

Table 2 Parameters and R^2 fitting to out of sample data

	parameters			R^2 factor model	
year	$a[MWh]$	$b[E]$	$c[€^{-1}]$	$I(t)$ only	$I(t)$ and $I_{tec}(t)$
2003	940.045	1.053	0.035	0.6948	0.7252
2004	901.082	1.113	0.039	0.4448	0.6170
2005	918.633	1.089	0.039	0.6069	0.7518
2006	902.485	1.072	0.038	0.8424	0.8499
2007	928.798	1.081	0.037	0.7571	0.7998
2008	884.110	1.068	0.043	0.6530	0.7025

the case of the gradient effect and the arbitrage incentive, respectively. This clearly

demonstrates the existence of non-random predictable components in the balancing energy demand data.
The raw data are obtained from the publications of balancing energy demand the four German TSOs from January 2003 to December 2008. Namely the four sources are RWE RWE Transportnetz Strom (2009), E.on e.on Netz (2009), EnBW EnBW Transportnetze (2009), and Vattenfall Vattenfall Europe Transmission (2009). These data from four sub-zones are combined to obtain one hypothetical German zone to account for interzonal balancing effects and provide a better correspondence to the single German day-ahead market EEX.

5 SARIMA model of balancing energy demand

Here we focus on the analysis of the residual time series after subtracting the effects described by the analytical model given by Equations (5) and (6). Before trying to adapt a linear time series model to the residuals, the data need to be checked for stationarity. We apply an augmented Dickey-Fuller unit root test. The null hypothesis of a unit root is rejected at a significance level below $\alpha = 0.001$ even when including the first 168 lags for the regression.
This finding is supported also by a consideration of the physical boundary conditions of the underlying data. Balancing energy demand is fulfilled by the grid operator to ensure grid balance. This energy has to be delivered physically by power stations, and so the installed capacity imposes a hard boundary. This boundary can, however, never be reached, as the response time and response capacity of power stations impose an even tighter boundary. Due to their design, powers stations cannot run on an arbitrary fraction of their designed capacity, but instead have to be operated within a certain bandwidth. Additionally, a complex system such as a power station has a considerable amount of inertia, and cannot instantaneously adapt to changes in operation. When looking at the total generation stock, although these facts do not translate into a hard boundary, the true limits will depend on the exact condition and history of all individual facilities connected to the grid. Nonetheless a limit to fluctuations the grid operator can manage always exists. So it is physically impossible for the balancing energy demand to grow to very large positive values or fall to very small negative values, but balancing energy will always be within a bandwidth around zero. Mathematically, this argument relates to a stationary time series, and the absence of unit roots. We can therefore model the data without the need of further differencing.
As a first step, an inspection of the SACF and SPACF of the residuals in Figure 5 shows the presence of SARIMA effects in the data. The autocorrelation decays off with increasing lag. Additionally, this decay is disturbed at multiples of 24, indicating a seasonality of 24. This picture is supported by the partial autocorrelation function displaying a drop at lag one and 24, together with a decaying negative partial autocorrelation at lags following multiples of 24. Moreover, the SPACF indicates another step at lag three. We therefore choose SARIMA$(1,0,0) \times (1,0,1)_{24}$ and

Fig. 5 Sample autocorrelation and partial autocorrelation

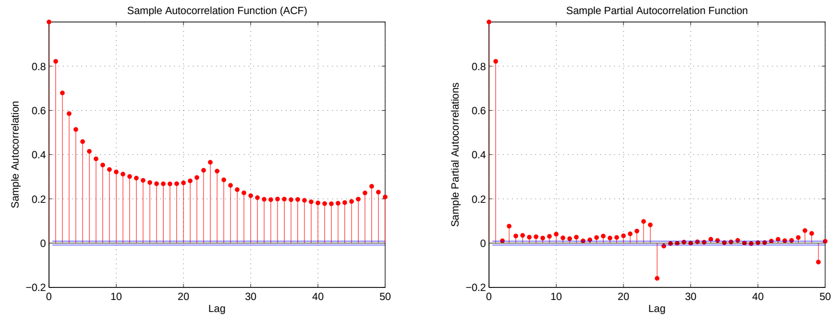

SARIMA$(3,0,0)\times(1,0,1)_{24}$ as candidates for the model. Additional to the classical model with Gaussian innovations we include models with t-distributed innovations as representatives of heavy-tailed innovations in the analysis. This approach is suggested by Zumbach (2006) as a compromise between a heavy-tailed innovation distribution and robust parameter estimates for the SARIMA-model. Table 3 holds the AIC and BIC values of different specifications, including both Gaussian and t-distributed innovations. We choose the SARIMA$(1,0,0)\times(1,0,1)_{24}$ model with t-distributed innovations for two reasons. First, the AIC and BIC values indicate a preference of t-distributed innovations over the Gaussian case. Second, the ar_3 coefficient is small and, as we will discuss below, coefficients at low lags are of minor practical relevance. Note that the t-distributed innovations demonstrate the necessity of a heavy-tailed noise term in the model. This will be further investigated in Section 6.

We conclude this section with an analysis of the consistency of the model over time and a test of its forecasts. Table 4 reports the parameter estimates of the model based

Table 3 Parameter estimates of SARIMA model

	SARIMA			
	$(1,0,0)\times(1,0,1)_{24}$		$(3,0,0)\times(1,0,1)_{24}$	
	Gaussian	t(ν)	Gaussian	t(ν)
ar_1	0.8185	0.8238	0.7974	0.8036
ar_3	-	-	0.0084	0.0079
ar_{24}	0.9572	0.9571	0.9419	0.9427
ma_{24}	-0.8502	-0.8529	-0.8291	-0.8328
σ	341.9658	341.9410	341.7744	342.0834
ν	-	9.1080	-	9.1396
AIC	763,170	762,080	763,060	761,970
BIC	763,200	762,110	763,090	762,000

on yearly sub-samples. The parameter estimates are very consistent with the overall model. We therefore decided to test the forecasts of the overall model rather than the individual yearly models.

The information on balancing energy demand is not continuously revealed to the market, but rather published only once a month, including the data for the preceding month. So the data for May will be available by July. We therefore do not test the one-time step forecast, as this has no practical implication in this market. Instead, we test a forecast adapted to the information revealed to the market. As a result, forecasting is performed once a month based on the information lagged one month (i.e. the forecast horizon is 720 to 1,440 lags). Because one of the TSOs also publishes the balancing demand data in its zone with a time delay of only three days, we also test the implications of the model on this time horizon (i.e. a forecast horizon of 72 to 96 lags). In both cases that we test, the sample variance is reduced by subtracting the conditional expectation. When using the monthly forecast, the variance is reduced by 3.69%; applying a three-day forecast horizon results in a 11.22% reduction.

Table 4 Parameter estimates of SARIMA model

parameter	a_1	a_{24}	b_{24}	σ	ν
total	0.8238	0.9571	-0.8529	341.9410	9.1080
2003	0.8522	0.9307	-0.7973	341.9410	9.4983
2004	0.7860	0.9421	-0.8184	328.8776	12.3724
2005	0.7810	0.9359	-0.8248	325.6431	10.7623
2006	0.8092	0.9455	-0.8397	336.8019	9.0538
2007	0.8067	0.9586	-0.8620	329.9785	10.1529
2008	0.7861	0.9515	-0.8541	341.9410	9.2676

This additional variance reduction as compared to the analytical model in Section 4 can be decomposed into two components. For the first component there exist comparatively short-lived patterns in the data. These patterns can be understood as a linear correction term for the analytical model. The second component captures a non-zero conditional mean of the time series. Neither the gradient effect nor the arbitrage incentive described in Section 4 can explain a non-zero expectation value of the balancing energy demand over extended periods of time. Both effects will average to zero over a few cycles of their respective seasonality. However when looking at the average forecast of the SARIMA-model in Table 5, it is evident the SARIMA forecast does not average to zero over a few cycles. Furthermore, it is argued in Möller et al (2009) that this non-zero mean over extended periods of time is sufficient to influence the electricity price on the day-ahead market. Other studies of market power abuse in Germany identify the same periods as periods of abuse, which show highly negative mean SARIMA forecasts (i.e., artificial demand) in their analysis (see EU (2007), von Hirschhausen et al (2007) and Schwarz and Lang (2006)).

Table 5 Average prediction of SARIMA model

horizon	average prediction 2003	2004	2005	2006	2007	2008
one month	-198.9537	-179.9403	-110.6992	21.8594	-118.4479	-52.1579
three days	-372.3728	-307.2270	-214.6126	18.6222	-217.6485	-190.4806

6 Innovation distribution

The parameter estimation of the SARIMA-model used in Section 5 is based on a t-distributed innovation process. Models with Gaussian innovations were disregarded based on AIC and BIC values reported in Section 5. In this section, we will further investigate the heavy-tailed innovation process.
The innovations process is of particular importance in the balancing energy market as it governs the risk involved with balancing the grid. In general, TSOs have to allocate sufficient regulatory power to be able to maintain grid operation and avoid a blackout. The capacity that is considered sufficient is usually defined by a threshold probability for a blackout (i.e., the probability of fluctuations exceeding the allocated capacity). So the more precise the quantiles of the innovations distributions are known, the more efficiently resources may be allocated.
The first step of the investigation is the QQ-plot of the innovations time series and the fitted t-distribution in Figure 6. From this figure it can be seen that the t-distribution does not adequately capture the risk in the tails of the empirical distribution because the QQ-plot deviates from the diagonal.
Due to the conceptual advantage of modeling the data with a distribution in the proximity of the GCLT as discussed in Section 3, we test both the α-stable and the CTS-distribution as a more adequate model for the innovations time series. Both distributions are estimated by the Fourier inversion formula and their characteristic functions given by Equations (3) and (4). This inversion is, in turn, numerically estimated by the fast Fourier transform (FFT) method. A more detailed description of the method is given in Nolan (1997) and Kim et al (2009). Table 6 shows the estimated parameter sets. As can be seen in Figure 7, the heavy-tailed distribution captures the likelihood of extreme events more accurately than the t-distribution.

Table 6 Estimated parameters of heavy-tailed distributions

distribution	parameters				
α-stable	α	σ	β		μ
	1.9107	0.0048	0.6711		0-fixed
CTS	α	C	λ_+	λ_-	m
	0.9122	$\frac{1}{\Gamma(2-\alpha)(\lambda_+^{\alpha-2}+\lambda_-^{\alpha-2})}$	1.4856	1.5168	0-fixed

Fig. 6 QQ-plot *t*-distribution

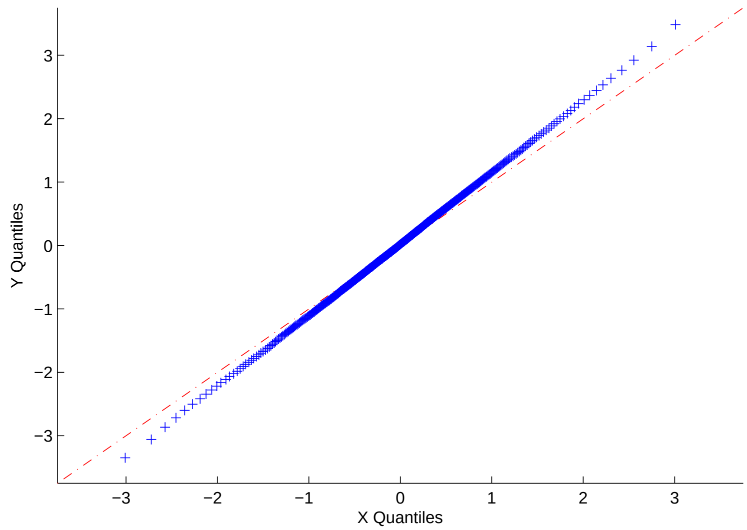

Fig. 7 QQ-plot CTS-distribution

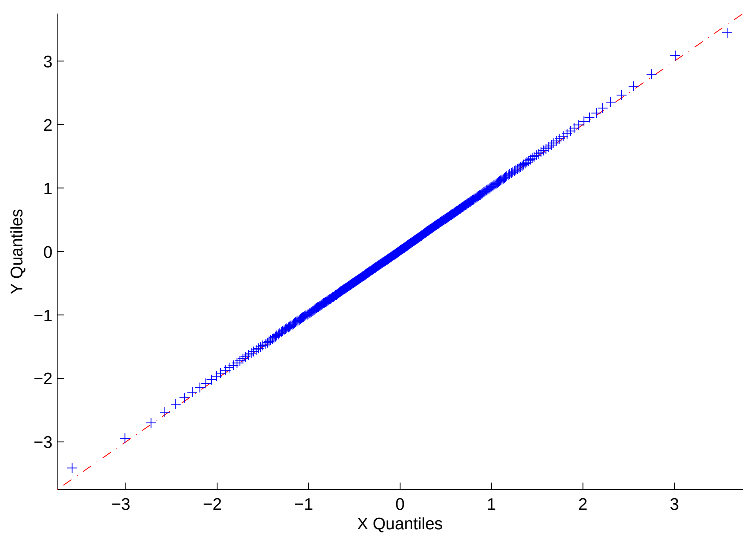

In the next step, all three distributions are compared using the goodness-of-fit tests mentioned in Section 3. The results are summarized in Table 7.
The p-values of the KS-test clearly indicate that the CTS-distribution describes the

Table 7 Goodness-of-fit statistics and p-values

	distribution					
	t		α-stable		CTS	
test	statistic	p-value	statistic	p-value	statistic	p-value
KS	0.0331	$1.37 \cdot 10^{-50}$	0.0156	$1.75 \cdot 10^{-11}$	0.0082	0.0016
AD	0.0883	-	0.0333	-	0.0170	-
KS*	0.0316	$5.73 \cdot 10^{-46}$	0.0134	$1.40 \cdot 10^{-8}$	0.0059	0.0538
AD*	0.0854	-	0.0316	-	0.0126	-
CvM*	23.4809	-	2.2939	-	0.5144	-
AD^2*	625.3888	-	685.0225	-	611.5575	-

innovations best, as its p-value is 49 and 9 orders of magnitude greater than the p-value of t-distribution and α-stable distribution, respectively. However the p-value of the CTS-distribution is still low. As discussed in Section 3, the KS-test is responsive to small fluctuations in the location parameter, while such fluctuations are to be expected with heavy-tailed distributions. Also, the SARIMA model implies a location parameter of zero for the innovation process, so we do not need to focus on the location parameter. We therefore correct the mean of the innovations time series for such fluctuations within the 95% confidence bounds. The corresponding statistics are identified by an asterisk (*). Again the CTS-distribution provides the best description of the data. Furthermore, the CTS-distribution is acceptable at a 5% significance level. The other statistics reported provide further support for selecting the CTS-distribution over both the t-distribution and the α-stable distribution.

7 Conclusion

The wide application of linear time series models in finance has made the shortcomings of inadequate innovation processes and correlation structures evident. Tempered stable distributions have been proposed and tested on financial time series data to overcome these problems.
The tempered stable distribution is an expansion of the α-stable distribution. It combines heavy-tailed innovations over multiple timescales with finite higher moments. These properties make tempered stable distributions an excellent choice for modeling phenomena dependant on extreme events, which are at the same time bounded by other considerations.

In this paper, we apply the classical tempered stable model to German balancing energy demand data and demonstrate its fitness. The CTS-distribution describes the risk of unpredictable events in the electricity grid, while at the same time capturing physical boundary conditions in the model. In the balancing energy market, these advantages can be used to allocate the resource capacity more efficiently, and thereby reduce grid tariffs.
Additionally, the SARMIA model is able to separate predictable balancing energy demand from unpredictable shocks. In general, the predictable fraction of balancing energy demand can be satisfied by a wider and technically less demanding range of providers. The German market design includes the option to trade this predictable fraction in strategic positions on the balancing market, while the reserve capacity for the unpredictable shocks is traded on the market for regulatory energy. A further development of strategic positions in the balancing energy market could therefore increase the competitiveness of the capacity reserve market. In view of the planned increase of highly fluctuating regenerative power in the German system, this advantage will become even more pronounced in the future.

References

Box G, Jenkins G (1970) Time Series Analysis: Forecasting and Control. Holden-Day, San Francisco

Boyarchenko SI, Levendorskiĭ SZ (2000) Option pricing for truncated Lévy processes. International Journal of Theoretical and Applied Finance 3:549–552

Bundesnetzagentur (2006) Bilanzkreisvertrag über die Führung eines Bilanzkreises. http://www.bundesnetzagentur.de

Carr P, Geman H, Madan D, Yor M (2002) The fine structure of asset returns: An empirical investigation. Journal of Business 75(2):305–332

Chernobai A, Rachev S, Fabozzi F (2007) Operational Risk. The Frank J. Fabozzi Series, Wiley Finance, Hoboken, New Jersey

Dickey D, Fuller W (1979) Distribution of the estimators for autoregressive time series with a unit root. Journal of the American Statistical Association 74(366):427–431

EnBW Transportnetze (2009) Bilanzkreisabrechnung: Verfahren und Preise. https://www.regelleistung.net

eon Netz (2009) Bilanzkreisabweichung. https://www.regelleistung.net

EU (2005) Technical Annexes to the Report from the Commission on the Implementation of the Gas and Electricity Internal Market COM(2004)863 final. Commission staff working document, Commission of the European Communities

EU (2007) Structure and Performance of Six European Wholesale Electricity Markets in 2003, 2004 and 2005. European commission dg competition report, London Economics and Global Energy Decisions

von Hirschhausen, Weigt, Zachmann (2007) Preisbildung und Marktmacht auf den Elektrizitätsmärkten in Deutschland. Expert opinion(VIK), Universität Dresden EE2

Kim YS, Rachev ST, Bianchi ML, Fabozzi FJ (2008) Financial market models with Lévy processes and time-varying volatility. Journal of Banking and Finance 32:1363–1378

Kim YS, Rachev ST, Bianchi ML, Fabozzi FJ (2009) Computing var and avar in infinitely divisible distributions, technical report, Chair of Econometrics, Statistics and Mathematical Finance School of Economics and Business Engineering University of Karlsruhe (http://www.statistik.uni-karlsruhe.de/download/AVaR_ID_KRBF_03-25.pdf)

Koponen I (1995) Analytic approach to the problem of convergence of truncated Lévy flights towards the gaussian stochastic process. Physical Review E 52:1197–1199

Menn C, Rachev ST (2008) Smoothly truncated stable distributions, GARCH-models, and option pricing. Mathematical Methods in Operation Research To appear

Menn C, Rachev ST, Fabozzi FJ (2005) Fat-Tailed and Skewed Asset Return Distributions: Implications for Risk Management, Portfolio selection, and Option Pricing. John Wiley & Sons

Möller C, Rachev S, Kim YS, Fabozzi F (2009) Strategic deployment of balancing energy in the German electricity market. Working paper

Nailis D, Ritzau M (2006) Studie zur Marktausgestalltung der Regel- Ausgleichsenergie vor dem Hintergrund des neuen EnWG. Tech. rep., BET

Nolan JP (1997) Numerical calculation of stable densities and distribution functions. Communications in Statistics - Stochastic Models 13:759–774

Rachev ST, Mittnik S (2000) Stable Paretian Models in Finance. John Wiley & Sons, New York

RWE Transportnetz Strom (2009) Bilanzkreisabweichungen. https://www.regelleistung.net

Schwarz HG, Lang C (2006) The Rise in German Wholesale Electricity Prices: Fundamental Factros, Eexercise of Market Power, or Both? Tech. rep., University of Erlangen-Nürnberg IWE

Vattenfall Europe Transmission (2009) Ausgleichsenergie Arbeitspreise. https://www.regelleistung.net

Zumbach G (2006) A gentle introduction to the RM 2006 methodology. Technical report, RiskMetrics Group