Discussion Papers in Economics

No. 17/2017

Optimal trading strategies for Lévy-driven Ornstein-Uhlenbeck processes

Sylvia Endres
University of Erlangen-Nürnberg

Johannes Stübinger
University of Erlangen-Nürnberg

ISSN 1867-6707

Friedrich-Alexander-Universität Erlangen-Nürnberg
Institute for Economics
https://www.iwf.rw.fau.de/research/iwf-discussion-paper-series/

Optimal trading strategies for Lévy-driven Ornstein–Uhlenbeck processes

Sylvia Endres[a,1], Johannes Stübinger[a,1]

[a] *University of Erlangen–Nürnberg, Department of Statistics and Econometrics, Lange Gasse 20, 90403 Nürnberg, Germany*

Monday 25th September, 2017

Abstract

This study derives an optimal pairs trading strategy based on a Lévy-driven Ornstein–Uhlenbeck process and applies it to high-frequency data of the S&P 500 constituents from 1998 to 2015. Our model provides optimal entry and exit signals by maximizing the expected return expressed in terms of the first-passage time of the spread process. An explicit representation of the strategy's objective function allows for direct optimization without Monte Carlo methods. Categorizing the data sample into 10 economic sectors, we depict both the performance of each sector and the efficiency of the strategy in general. Results from empirical back-testing show strong support for the profitability of the model with returns after transaction costs ranging from 31.90 percent p.a. for the sector "Consumer Staples" to 278.61 percent p.a. for the sector "Financials". We find that the remarkable returns across all economic sectors are strongly driven by model parameters and sector size. Jump intensity decreases over time with strong outliers in times of high market turmoils. The value-add of our Lévy-based model is demonstrated by benchmarking it with quantitative strategies based on Brownian motion-driven processes.

Keywords: Finance; Pairs trading; Optimal thresholds; Ornstein–Uhlenbeck Lévy process; Mean-reversion; High-frequency data

Email addresses: sylvia.endres@fau.de (Sylvia Endres), johannes.stuebinger@fau.de (Johannes Stübinger)

[1] The authors have benefited from many helpful discussions with Ingo Klein, Julian Knoll, and Christopher Krauß.

1. Introduction

Pairs trading is a market neutral strategy which has been developed by a group of computer scientists, mathematicians, and physicists at Morgan Stanley in the early to mid-1980s (Vidyamurthy, 2004). In its most common form, this strategy finds pairs of related stocks whose prices follow a similar pattern over a historical time period. In case of abnormal divergence, an arbitrageur goes long in the undervalued stock while simultaneously going short in the overvalued stock. The exposure to market risk is substantially reduced for this kind of trading strategy. If history repeats itself, the price relationship converges to the long-term equilibrium and a profit is taken. Gatev et al. (2006) expound the first major academic study on statistical arbitrage pairs trading exhibiting excess returns of 11 percent p.a. for U.S. CRSP securities from 1962 until 2002. This seminal paper characterizes the trigger for an ever-expanding interest in this field of research up to the present. Substantial contributions are supplied by Vidyamurthy (2004), Elliott et al. (2005), Avellaneda and Lee (2010), Do and Faff (2012), Huck and Afawubo (2015), and Liu et al. (2017).

The vast majority of the literature determines model-free trading rules via a pre-defined parameter setting, e.g., positions are opened at a two-standard deviation spread — the criticism of data snooping is omnipresent. Only a small fraction depicts analytic formulae and solutions for calculating optimal model-driven statistical arbitrage pairs trading thresholds. This research is confined to Elliott et al. (2005), Bertram (2010a,b), Ekström et al. (2011), Govender (2011), Gregory et al. (2011), Cummins and Bucca (2012), Zeng and Lee (2014), Leung and Li (2015a,b), Li (2015), Göncü and Akyıldırım (2016a), Bai and Wu (2017), Baviera and Baldi (2017), and Suzuki (2017). These studies model the price spread between two stocks with an Ornstein–Uhlenbeck (OU) process based on Brownian motion.

Elliott et al. (2005) supply a seminal cornerstone by specifying the spread of two stocks via a mean-reverting Gaussian Markov chain model which is observed in Gaussian noise. Predictions from the calibrated model are used to take appropriate investment decisions. Bertram (2010a) derives analytic formulae for optimal statistical arbitrage trading assuming an OU process. Particularly, solutions for the optimal entry and exit thresholds are determined by maximizing the expected return per unit time. Bertram (2010b) presents ideal trading strategies for a two-factor model with zero drift and independence between the two associated Brownian motions. In an empirical back-testing study, Cummins and Bucca (2012) apply the model of Bertram (2010a) to the crude oil and refined products markets from 2003 to 2010. Govender (2011), Gregory et al. (2011), Li (2015), and Leung and Li (2015a) occupy oneself with the optimal timing of trades supposing different underlying frameworks, i.e., OU model, exponential OU model, Cox–Ingersoll–Ross model, and GARCH model. These strategies are applied to different stock markets and varying time periods. Ekström et al. (2011) analyze the optimal single liquidation timing and consider the ideal trade liquidation regarding a discount factor. Zeng and Lee (2014) find optimal thresholds using a polynomial expression for the expectation of the first-passage time of an OU process with two-sided boundary. Ideal entry and exit decision rules for a double stopping problem are analytically shown by Leung and Li (2015b). Göncü and Akyıldırım (2016a) present a new filtering technique to identify the most suitable pairs in the market. Bai and Wu (2017) analyze optimal stopping for Markov-modulated OU processes. Baviera and Baldi (2017) derive analytically the optimal trading thresholds as well as the ideal leverage for every given stop-loss level and Suzuki (2017) considers the optimal multiple switching problem for OU processes.

Given the characteristics of financial data, OU processes based on Brownian motion are not eligible to model high-frequency dynamics — direct consequences are invalid parameter estimates and disregard of stylized facts (Barlow 2002, Carr et al. 2002, Cont and Tankov 2003, Cartea and Figueroa 2005, Meyer-Brandis and Tankov 2008, Jing et al. 2012, Aït-Sahalia and Jacod 2014, Jondeau et al. 2015, Göncü and Akyıldırım 2016b). Surprisingly, only Larsson et al. (2013) and Göncü and Akyıldırım (2016b) present academic studies in the context of determining optimal thresholds using Lévy-driven OU processes. In this respect it is important to note that both exhibit merely numerical solutions. Larsson et al. (2013) generalize the approach of Ekström et al. (2011) by additionally including a jump term. Then, the authors prove a verification theorem and evaluate a numerical method for the corresponding free boundary problem. Göncü and Akyıldırım (2016b) give attention to Lévy-driven OU processes with generalized hyperbolic distributed marginals. Together, these two studies provide an initial step for determining optimal trading strategies based on Lévy-driven OU processes.

We enhance the existing research in several aspects. First, our manuscript contributes to the literature by developing an optimal trading framework based on Lévy-driven OU processes, i.e., we derive optimal trade thresholds by maximizing the expected return per unit time. In stark contrast to the existing literature, Monte Carlo methods are avoided due to a mathematical expression for the expected first-passage time and an explicit representation of the strategy's objective function. Besides mean-reversion, volatility clusters, and drifts, this general and flexible class of stochastic models is able to capture jumps and fat tails. Considering all these effects, our strategy is able to perform both intraday and overnight trading. Most notable, we generalize the Brownian motion-driven OU processes by including the essential jump component. As such, our model is in a position to choose the borderline case in times of full stagnation or quiet behavior without fat tails in returns at the stock market. Second, we present the first academic survey applying a high-frequency back-testing study to all industry sectors of the S&P 500 over a sample period of 18 years. Therefore, we categorize all companies into 10 economic sectors based on the Global Industry Classification Standard (GICS) and consider all pairs consisting of stocks from the same sector. Thus, we are in a position both to demonstrate the performance of each sector and the efficiency of our strategy in general. For all sectors, we observe statistically and economically significant returns after transaction costs from 1998 to 2015, ranging from 31.90 percent p.a. for the sector "Consumer Staples" to 278.61 percent p.a. for the sector "Financials". Third, we analyze common effects across all economic sectors and observe that returns are strongly driven by model parameters and sector size. This finding explains the varying performance between the different branches of trade. Forth, we demonstrate the value-add of our optimal trading strategy for Lévy-driven OU processes by benchmarking it with well known quantitative strategies in the same field. Specifically, we consider a classic asymmetric strategy in the sense of Zeng and Lee (2014) and an optimal trading strategy by Bertram (2010a), both based on a Brownian motion-driven OU process. We find that our strategy outperforms the classic pairs trading approaches over time and obtains almost no loading on systematic sources of risk.

The rest of this paper is organized as follows. Section 2 describes the data set and software utilized in this study. In section 3, we depict the theoretical framework for determining optimal trading strategies. Section 4 provides the study design for our back-testing application. Results and key findings are reported and discussed in section 5. Finally, section 6 summarizes the main results of this paper and makes methodological recommendations about further research.

2. Data and Software

We run our back-testing framework on minute–by–minute prices of the S&P 500 index constituents from January 1998 until December 2015. Our highly liquid trading universe comprises the 500 leading companies which satisfy strict requirements based on market size, liquidity, and industry grouping. Focusing on large-cap U.S. equities, this subset covers 80 percent of available market capitalization (S&P 500 Dow Jones Indices, 2015). Consequently, this market segment poses as a crucial test for any potential stock market inefficiency. Following Stübinger and Endres (2017), we eliminate the survivor bias from our data with the aid of a 2-stage process. First, using QuantQuote (2016), we produce a daily constituent list for the S&P 500 stocks from 1998 to 2015. Then, exploiting these information, we create a binary matrix, i.e., for each element of this matrix, we assign a "1" if the corresponding company is a constituent of the S&P 500 index at the current day, otherwise a "0". Second, for all companies having ever been part of the S&P 500 index, we download the minute–by–minute stock prices from January 1998 until December 2015 from QuantQuote (2016). An adjustment of the data set is conducted due to stock splits, dividends, and further corporate actions. Applying these two steps, we are in a position to completely replicate the S&P 500 index constituency and the respective prices over time.

The presented methodology in this paper and all relevant analyses are implemented in the statistical programming language `R` (R Core Team, 2017). For computationally intensive tasks the majority of source code is implemented in `C++` and linked to `R`. For time series handling, we rely on the packages `TTR` by Ulrich (2016) and `xts` by Ryan and Ulrich (2014). The vast majority of the performance evaluation is employed via the package `PerformanceAnalytics` by Peterson and Carl (2014).

3. Methodology

According to Zeng and Lee (2014), the concept of finding an optimal pairs trading strategy follows a three-step process. First, we specify a stochastic model explaining the spread dynamics. In this study, we assume a Lévy-driven Ornstein–Uhlenbeck (OU) process — a highly flexible approach for capturing typical characteristics of financial data, among them jumps, mean-reversion, volatility cluster, and drifts. Most notable, our model also includes the Brownian motion-driven OU process. In the second step, we define the target function of the strategy, e.g., maximize the expected profit or minimize the time until trade termination. Finally, we aim for deriving entry and exit times by optimizing the target function. This section describes the 3-step logic outlined above in detail.

3.1. Stochastic model

We define the spread for two stocks A and B with prices $S_A(t)$ and $S_B(t)$ as

$$X_t = \ln(S_A(t)/S_A(0)) - \ln(S_B(t)/S_B(0)),\ t \geq 0.$$

For modeling the spread dynamics, we consider the following OU process driven by Lévy noise:

$$dX_t = \theta(\mu_t - X_t)dt + dL_t,\ X_0 = x, \tag{1}$$

with mean-reversion speed $\theta \in \mathbb{R}$, time-dependent mean-reversion level $\mu_t \in \mathbb{R}$, and Lévy process $\{L_t\}_{t\geq 0}$. As a special case, the model covers earlier spread models based on Brownian motion-driven OU processes (Elliott et al. 2005, Bertram 2009, Bertram 2010a, Ekström et al. 2011, Cummins and Bucca 2012, Bogomolov 2013, Göncü and Akyıldırım 2016a). The solution of equation (1) is described by

$$X_t = xe^{-\theta t} + \theta \int_0^t \mu(u)e^{-\theta(t-u)}du + \int_0^t e^{-\theta(t-u)}dL_u.$$

In line with Kou (2002), Kou and Wang (2003), Sepp (2003), Ramezani and Zeng (2007), Bayraktar and Xing (2011), Cai and Kou (2011), Gerber et al. (2013), Kou et al. (2017), and Su and Bai (2017), we assume a double exponential jump–diffusion representation for the Lévy process $\{L_t\}_{t\geq 0}$, which fits stock data better than a model with normally distributed jump sizes (Ramezani and Zeng 1998, Li et al. 2008, Kou et al. 2017).

The sequence of random variables is divided into two parts — a continuous part driven by a Brownian motion and a compound Poisson process possessing a double exponential distribution. Mathematically, we define the Lévy process by

$$L_t = \sigma W_t + \sum_{i=1}^{N_t} \xi_i, \tag{2}$$

with standard Brownian motion $\{W_t\}_{t\geq 0}$, volatility $\sigma \in \mathbb{R}^+$, Poisson process $\{N_t\}_{t\geq 0}$ with rate λ, and jump sizes $\{\xi_1, \xi_2, \dots\}$. The processes $\{W_t\}_{t\geq 0}$, $\{N_t\}_{t\geq 0}$, and the random variables $\{\xi_1, \xi_2, ...\}$ are independent. The common density of ξ is given by

$$f_\xi(x) = p_u \eta_u e^{-\eta_u(x-\delta)} \mathbb{1}_{\{x\geq\delta\}} + p_d \eta_d e^{\eta_d(x+\delta)} \mathbb{1}_{\{x\leq-\delta\}},$$

where $p_u, p_d \in \mathbb{R}_0^+$ with $p_u + p_d = 1$, $\eta_u, \eta_d \in \mathbb{R}^+$, and $\delta \in \mathbb{R}_0^+$. The constants p_u and p_d represent the probabilities of upward (u) and downward (d) jumps. The indicator functions only admit jump sizes greater (smaller) than δ ($-\delta$). The moment generating function of the jump size ξ is given by

$$E[e^{s\xi}] = p_u \frac{\eta_u}{\eta_u - s} e^{\delta s} + p_d \frac{\eta_d}{\eta_d + s} e^{-\delta s}, \quad s \in (-\eta_d, \eta_u),$$

from which the moment generating function of the process $\{L_t\}_{t\geq 0}$ is obtained as

$$E[e^{sL_t}] = e^{\lambda t (p_u \frac{\eta_u}{\eta_u - s} e^{\delta s} + p_d \frac{\eta_d}{\eta_d + s} e^{-\delta s} - 1)} e^{\sigma^2 s^2 t/2}. \tag{3}$$

In the double exponential jump–diffusion model, the upward and downward jumps are both generated by a single Poisson process with fixed intensity λ. Ramezani and Zeng (1998) propose a similar model, in which upward and downward jumps are caused by two independent Poisson processes $\{N_t^j(\lambda^j)\}_{t\geq 0}$ with intensity parameters λ^j ($j \in \{u, d\}$). The corresponding Lévy process is defined by

$$L_t = \sigma W_t + \sum_{j\in\{u,d\}} \sum_{i=1}^{N_t^j(\lambda^j)} \xi_i^j. \tag{4}$$

The density functions for jump magnitudes are $f_{\xi^u}(x) = \eta_u e^{-\eta_u(x-\delta)} \mathbb{1}_{\{x\geq\delta\}}$ and $f_{\xi^d}(x) = \eta_d e^{\eta_d(x+\delta)} \mathbb{1}_{\{x\leq-\delta\}}$. Besides economic reasons, the distinction between upward and downward variations leads to better analytical tractability. Following Ramezani and Zeng (2007), we are in a position to transform the parameters of (2) and (4) into each other by setting $\lambda = \lambda_u + \lambda_d$ and $p_u = \frac{\lambda_u}{\lambda}$. Furthermore, we set $\eta_u = \eta_d = \eta$ and $p_u = 1/2$ (Carr and Wu 2004, Tsay 2005, Ramezani and Zeng 2007, Huang and Huang 2012) — an assumption that has been commonly invoked.

The proposed Lévy framework is able to reproduce the leptokurtic feature of the return distribution, with a higher peak and heavier tails than those of the normal distribution — this links well with empirical features of the spread process that are clearly non-Gaussian (Bertram 2009, Aït-Sahalia and Jacod 2014, Göncü and Akyıldırım 2016b).

3.2. Objective function

As a naturally target of any rational investor, we maximize the expected return per unit time in our trading strategy. We define the entry and exit signals x and b ($x, b \in \mathbb{R}$) and assume $x < b$ without loss of generality. The time over which the return takes place is the first-passage time $\tau_{b,x}$:

$$\tau_{b,x} = \inf\{t \geq 0 | X_t \geq b\}$$

for $X_0 = x$. Figure 1 illustrates the setting for simulated spread dynamics and thresholds x and b. Positions are entered and exited when the spread crosses the thresholds. The squares representing the respective times are not exactly on the dashed lines because the spread exhibits jumps and additionally, trading is discrete. The investor enters at $t_0 = \inf\{t \geq 0 | X_t \leq x\}$ and exits at $\tau_{b,x} = \inf\{t \geq 0 | X_t \geq b\}$ for the first time.

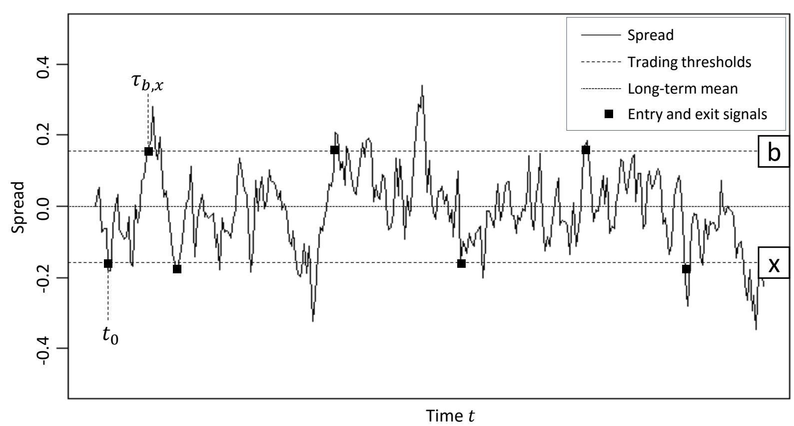

Figure 1: Trading strategy for simulated spread dynamics (black curve) and thresholds x and b (dashed lines). Squares label the times of opening and closing positions.

The return r is a function of the thresholds x and b and transaction costs c ($c \in \mathbb{R}_0^+$). This results in the optimization problem

$$\max_{b,x} \frac{r(b,x,c)}{E[\tau_{b,x}]}, \tag{5}$$

which is actually a problem of determining or approximating the distribution of the first-passage time $\tau_{b,x}$. In contrast to the Brownian motion-driven OU process, the first-passage time problem for general Lévy processes has not been solved analytically. Nevertheless, the construction of the Lévy process in equations (2) and (4) offers a rare case allowing for analytical solutions for the first passage times, derived in Borovkov and Novikov (2008). These involve an explicit representation of the mean of $\tau_{b,x}$.

3.3. Optimal variables

In this subsection, we determine the optimal deviation thresholds to generate trade signals and execute pairs trading. Therefore, we decompose the Lévy process $\{L_t\}_{t\geq 0}$ into two parts — $\{Q_t\}_{t\geq 0}$, representing the downward jumps and the Brownian motion, and $\{R_t\}_{t\geq 0}$, representing the upward jumps. Both jump sizes follow exponential distributions. Due to the memory-less property of the exponential distribution it holds that

$$E(\tau_{b,x}) = \frac{1}{\theta}\int_0^K (e^{ub} - e^{ux}(1-u/K))(1-u/K)^{\lambda/(2\theta)-1} e^{-\Delta(u)} u^{-1} du \tag{6}$$

for

$$\Delta(u) = \frac{1}{\theta}\int_0^u \frac{\log(Ee^{vQ_1})}{v} dv \tag{7}$$

and

$$K = \sup\{u \geq 0 | Ee^{uL_1} < \infty\} \tag{8}$$

with $X_0 = x$ (Borovkov and Novikov 2008). The moment generating functions in equations (7) and (8) are given by

$$E[e^{sQ_t}] = e^{\lambda_d t\left(\frac{\eta}{\eta+s}e^{-\delta s}-1\right)} e^{\sigma^2 s^2 t/2}$$

and

$$E[e^{sL_t}] = m_{Q_t}(s) e^{\lambda_u t\left(\frac{\eta}{\eta-s}e^{\delta s}-1\right)}.$$

Due to the explicit representation of $E(\tau_{b,x})$, we are in position to directly solve problem (5). Typically, such optimization problems are solved by a box-constrained optimization routine using the Port library (Gay 1990, Fox 1997). We obtain thresholds that maximize the expected return per unit time in the respective trading strategy.

4. Study design

For our back-testing framework, we follow Jegadeesh and Titman (1993) and Gatev et al. (1999, 2006) and divide our high-frequency data set from January 1998 to December 2015 into 4484 overlapping study periods (figure 2). Each study period is shifted by one day and consists of a 30-day formation period (subsection 4.1) and an out-of-sample 5-day trading period (subsection 4.2). We set the length of the formation and trading period consistent with Knoll et al. (2017), Liu et al. (2017), and Stübinger and Endres (2017).

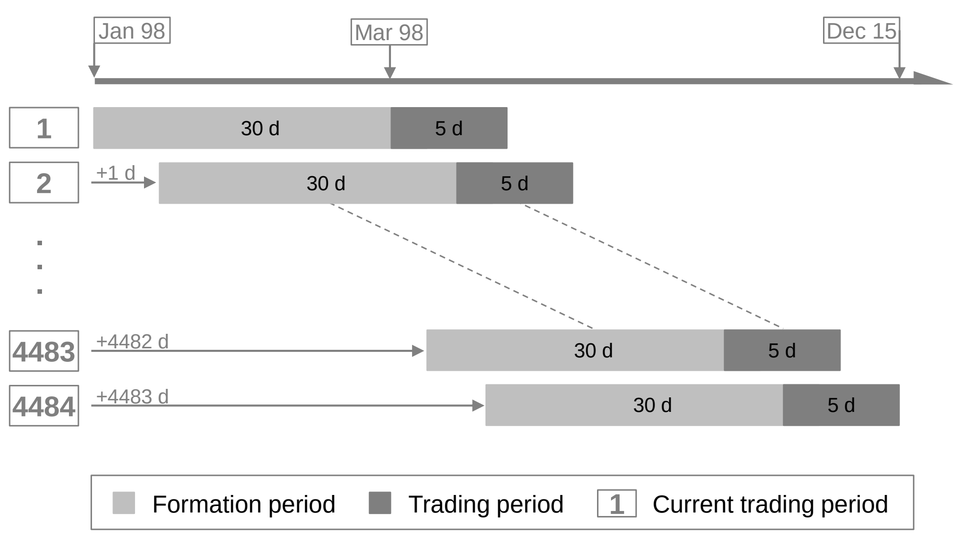

Figure 2: The back-testing application deals with 4484 overlapping study periods from January 1998 to December 2015. Each study period consists of a 30-day formation and a 5-day out-of-sample trading period.

4.1. Formation period

In the 30-day formation period, we fit Ornstein–Uhlenbeck (OU) processes driven by Lévy noise as specified in equation (1) to all possible combinations of pairs. Following Liu et al. (2017), we receive the mean-reversion level μ_t by a step function obtained from the average of the last two available data points out of daily opening and closing values. According to Mai (2012, 2014), there is a maximum likelihood estimation for the mean-reversion speed θ available. Given $X_{t_1}, ..., X_{t_n}$, the estimator is obtained from the discretization of the stochastic differential equation for $\Delta_i X = X_{t_{i+1}} - X_{t_i}$ and $\Delta_n = \max\limits_{1 \leq i \leq n-1} \{|t_{i+1} - t_i|\}$ under the assumption that the jumps are of finite activity:

$$\hat{\theta}_n = \frac{\sum_{i=0}^{n-1} (\mu_{t_i} - X_{t_i}) \Delta_i X \mathbb{1}_{\{|\Delta_i X| \leq \nu_n\}}}{\sum_{i=0}^{n-1} (\mu_{t_i} - X_{t_i})^2 (t_{i+1} - t_i)}.$$

Increments larger than the threshold $\nu_n = \Delta_n^{\beta}$, $\beta \in (0, 1/2)$, are deleted to approximate the continuous part with the remaining series. This is due to the fact that increments of the continuous part over an interval length of Δ_n are, with high probability, smaller than $\Delta_n^{0.5}$ (Mancini 2009, Mai 2012, Aït-Sahalia and Jacod 2014). The time interval Δ_n in our context is $\Delta_n = \frac{1}{250 \cdot 391}$ (Cont and Mancini 2011, Liu et al. 2017). We set the most restrictive variant $\beta = 0.4999$ to specify the threshold exponent, a value well in line with Cont and Mancini (2011), who use $\beta = 0.999/2$ in an application to the S&P 500 index. It should be stressed that the estimator does not detect any jumps in the special case of a Brownian motion-driven OU process (Mai 2012). The procedure in which the jump process is distinguished from the continuous part enables to separately estimate the parameter set of the jumps (λ, η) and the continuous part (θ, σ, μ_t). In accordance with Cartea and Figueroa (2005), the volatility σ is estimated by the sample standard deviation. The remaining jump parameters λ and η are calibrated by their maximum likelihood estimators.

We transfer the top 10 pairs (Miao 2014, Stübinger et al. 2016) exhibiting both a high mean-reversion speed θ and a high volatility σ to the trading period. A larger θ leads to a higher trading frequency, and a larger σ leads to a bigger fluctuation of the process, both resulting in a higher profit in each trade (Zeng and Lee 2014). This selection criterion is in line with Liu et al. (2017), who select pairs with low long-term and high short-term variance, both functions of θ and σ, and Stübinger and Endres (2017), who select pairs with high mean-reversion speed θ and high jump intensity λ.

4.2. Trading period

When constructing trading bands of the pairs trading strategy, literature distinguishes between two approaches — the asymmetric and the symmetric approach (see Zeng and Lee

2014). In the asymmetric strategy, which is the common one among practitioners, mean-reverting processes are traded using asymmetric bands. Positions are opened at some extreme value, e.g., a two-standard deviation event (Gatev et al. 2006), and closed when the spread reaches the long-term mean. The symmetric approach on the other hand constructs trading bands symmetric about the mean level of the process. Representatives of this stream are Bertram (2010a), Cummins and Bucca (2012), and Zeng and Lee (2014). Bertram (2010a) proves that for optimal trading in the OU model, the trade entry and exit levels are symmetric about the mean of the traded security. This is in line with Zeng and Lee (2014), who use a larger trade cycle where the trader holds the spread position until a larger profit level is attained. They prove that for the OU process and a given value of transaction costs c, the return for entry and exit levels b and $-b$ symmetric about zero is higher than for an exit level at 0 and an entry at b. This observation may be explained by the impact of transaction costs — in the symmetric approach, the transaction costs are halved compared to the asymmetric strategy. Cummins and Bucca (2012) use symmetric trading signals b and x with alternating interpretation between entry and exit signals. They go long (short) when the trading signal b (x) is breached and then close position once x (b) is crossed. Consequently, they aggregate two trading approaches, one taking long positions and the second taking short positions.

We stick to the symmetric approach and choose thresholds $\mu_t \pm b$ around the mean level μ_t of our process, setting $x = -b$. Following Bertram (2010a), the return per trade is defined by the range between upper and lower trading threshold, adjusted for transaction costs c ($c \in \mathbb{R}_0^+$). Hence, we receive a return per trade of $r(b, -b, c) = 2b - c$ and require $b \geq \frac{c}{2}$ for any profitable strategy. The optimal value for b is achieved as result of the optimization routine outlined in subsection 3.3.

From the symmetric construction and the relationship $\tau_{b,-b} = \tau_{-b,b}$, we are able to use both thresholds as entry and exit signals at the same time. From equation (5), we obtain

$$\max_{b,x} \frac{r(b, x, c)}{E[\tau_{b,x}]} = \max_{b} \frac{2b - c}{E[\tau_{b,-b}]}, \tag{9}$$

with $X_0 = x = -b$. We define the following trading rules:

- Go short in stock A and go long in stock B, if $X_t \geq \mu_t + b$, i.e., stock A is overvalued and stock B undervalued.

- Go long in stock A and go short in stock B, if $X_t \leq \mu_t - b$, i.e., stock A is undervalued and stock B overvalued.

- Do not execute any trade, if $\mu_t - b < X_t < \mu_t + b$, i.e., the spread is in its 'normal' region.

We make a trade immediately upon every entry signal, buying 1 dollar worth of the undervalued stock and shorting 1 dollar worth of the overvalued stock. Allowing for only one active position per pair, we neglect any additional entry signals until the position is closed. Trades are held until the spread crosses the opposite trading band, the trading period ends, or one of the stocks of the respective pair is delisted.

We follow Gatev et al. (2006) for return computation. The sum of daily payoffs across all pairs is related to the sum of invested capital at the beginning of the respective day. We show both the return on committed capital (scale payoffs by the number of considered pairs) and the return on employed capital (scale payoffs by the number of active pairs). A pair is called active if it possesses at least one round-trip trade during the corresponding trading period. Specifically, the return on committed capital at day t is determined using the following definition:

$$\text{Return at day } t = \frac{\text{Sum of net profits at day } t}{\text{Number of considered pairs at day } t}.$$

Similarly, the return on employed capital at day t is calculated as:

$$\text{Return at day } t = \frac{\text{Sum of net profits at day } t}{\text{Number of active pairs at day } t}.$$

For our back-testing application, we decide on all S&P 500 constituents from January 1998 to December 2015. According to the Global Industry Classification Standard (GICS), all companies are categorized into the following 10 economic sectors (valuation date: 2015/12/31): Consumer Staples (CS), Consumer Discretionary (CD), Energy (EN), Financials (FI), Health Care (HC), Industrials (IN), Information Technology (IT), Materials (MA), Telecommunications Services (TS), and Utilities (UT). We generate all pairs consisting of stocks from the same sector. By this procedure, we are in a position both to analyze the performance of each sector and to examine the efficiency of our strategy in general.

5. Results

In the following section, we run a full-fledged performance evaluation after transaction costs for the top 10 stocks of each sector from March 1998 to December 2015, compared to a naive buy–and–hold strategy of the S&P 500 index (MKT). Therefore, we examine the risk–return characteristics and trading statistics (subsection 5.1) and evaluate the strategy performance across the sectors (subsection 5.2). We analyze the profitability in recent years (subsection 5.3), and investigate the influence of jumps to financial and statistical factors (subsection 5.4). Finally, we demonstrate the value-add of our optimal trading strategy based on Lévy-driven Ornstein–Uhlenbeck (OU) processes to less complex approaches built on regular OU processes (subsection 5.5). In line with the vast majority of the literature, returns are calculated based on committed capital. We follow Avellaneda and Lee (2010) and Liu et al. (2017) and depict transaction costs of 5 bps per share per half-turn.

5.1. Risk–return characteristics and trading statistics

Table 1 reports daily risk–return characteristics after transaction costs for the top 10 pairs per sector from March 1998 until December 2015. Most of the presented performance metrics are examined by Bacon (2008). Across all branches of trade, we observe statistically significant returns with Newey–West t-statistics between 5.84 for CS and 16.00 for EN. The economic point of view confirms this finding — daily returns after transaction costs range from 0.12 percent for CS to 0.54 percent for FI, compared to 0.01 percent for the S&P 500 benchmark. In most instances the return distribution of the sectors possesses right skewness — a desired property for any investor (Cont, 2001). Kurtosis well above 3 indicates leptokurtic distribution for all variants. Following Mina and Xiao (2001), we depict historical Value at Risk (VaR) measures, typically used to gauge the amount of assets required to cover potential losses. Tail risks of all sectors are approximately at the same level as the market, e.g., the historical VaR 5% vary from -2.53 percent for IT to -1.13 percent for UT, compared to -1.97 percent for a naive buy–and–hold strategy. Most notable, the maximum drawdown level varies widely within the different branches of trade from 0.23 for IN to 0.76 for CS. Across all sectors, the hit rate, i.e., the percentage of days with positive returns, exceeds clearly the market (53.06 percent) with a top value of 68.02 percent for FI.

	CS	CD	EN	FI	HC	IN	IT	MA	TS	UT	MKT
Mean return	0.0012	0.0023	0.0040	0.0054	0.0024	0.0028	0.0040	0.0034	0.0023	0.0032	0.0001
Standard error (NW)	0.0002	0.0003	0.0002	0.0004	0.0003	0.0002	0.0004	0.0003	0.0003	0.0002	0.0002
t-Statistic (NW)	5.8389	7.6116	15.9977	14.7785	9.1559	11.9127	11.0501	12.3458	7.2136	13.9580	0.8889
Minimum	-0.1354	-0.1603	-0.1406	-0.1340	-0.2814	-0.1267	-0.1303	-0.0954	-0.2042	-0.1498	-0.0947
Quartile 1	-0.0037	-0.0048	-0.0025	-0.0016	-0.0045	-0.0036	-0.0053	-0.0036	-0.0056	-0.0017	-0.0056
Median	0.0009	0.0022	0.0033	0.0035	0.0024	0.0025	0.0034	0.0027	0.0017	0.0018	0.0005
Quartile 3	0.0056	0.0095	0.0100	0.0101	0.0095	0.0090	0.0129	0.0098	0.0097	0.0074	0.0061
Maximum	0.1580	0.1836	0.1552	0.1742	0.1020	0.1577	0.1711	0.1390	0.1641	0.1083	0.1096
Standard deviation	0.0110	0.0173	0.0125	0.0166	0.0154	0.0135	0.0201	0.0135	0.0184	0.0120	0.0126
Skewness	0.9116	0.2388	0.2152	1.2468	-1.8829	0.4082	0.4566	0.5782	-0.0652	-0.7257	-0.1983
Kurtosis	23.2571	14.9316	13.3305	15.4023	34.0079	10.3881	7.5709	7.0291	13.1650	21.4208	7.5250
Historical VaR 1 %	-0.0270	-0.0458	-0.0259	-0.0344	-0.0400	-0.0337	-0.0528	-0.0332	-0.0451	-0.0262	-0.0350
Historical CVaR 1 %	-0.0386	-0.0709	-0.0410	-0.0587	-0.0652	-0.0473	-0.0712	-0.0432	-0.0683	-0.0494	-0.0506
Historical VaR 5 %	-0.0145	-0.0219	-0.0136	-0.0120	-0.0186	-0.0170	-0.0253	-0.0164	-0.0234	-0.0113	-0.0197
Historical CVaR 5 %	-0.0229	-0.0382	-0.0224	-0.0268	-0.0337	-0.0277	-0.0416	-0.0266	-0.0388	-0.0234	-0.0302
Maximum drawdown	0.7575	0.4322	0.6629	0.3267	0.4838	0.2315	0.6944	0.7142	0.5174	0.3369	0.6433
Share with return > 0	0.5624	0.5883	0.6512	0.6802	0.5892	0.6115	0.6079	0.6200	0.5662	0.6363	0.5306

Table 1: Daily return characteristics and risk metrics after transaction costs for the top 10 pairs of the sectors CS, CD, EN, FI, HC, IN, IT, MA, TS, UT, compared to a S&P 500 long-only benchmark (MKT) from March 1998 until December 2015. NW denotes Newey–West standard errors with five-lag correction and CVaR the Conditional Value at Risk.

Table 2 contains statistics on trading frequency, which show a similar picture for all sectors. At least 9.98 of the 10 regarded pairs possess trading activity during the 5-day trading period. This high number is explained by the second part of our pair selection algorithm outlined in subsection 4.1 — high volatility of the spreads creates superior trading opportunities. The average time pairs are open ranges between 0.74 days for IT and 1.27 days for UT, 8 out of 10 sectors exhibit intraday trade durations. This short-term horizon of our strategy is based on the first part of our selection — high mean-reversion speed leads to a fast reinvestment of capital and diminishes the risk of losing financial resources caused by a divergent pair.

	CS	CD	EN	FI	HC	IN	IT	MA	TS	UT
Avg. traded pairs per 5-days period	9.9996	10.0000	10.0000	9.9989	10.0000	10.0000	10.0000	10.0000	9.9853	9.9770
Avg. round-trip trades per pair	5.6809	7.6069	7.0411	7.1301	7.3590	6.9598	8.6773	7.0300	7.4031	5.4587
St. dev. round-trip trades per pair	2.9196	3.5827	3.2275	4.0860	3.3319	3.2904	4.3661	3.3780	4.1741	3.3399
Avg. time pairs are open in days	1.1200	0.8209	0.8772	0.9459	0.8334	0.8951	0.7394	0.8973	0.9242	1.2670
St. dev. of time open, per pair, in days	0.6892	0.4737	0.5076	0.6417	0.4589	0.5137	0.4456	0.5424	0.6670	0.9070

Table 2: Trading statistics for the top 10 pairs of the sectors CS, CD, EN, FI, HC, IN, IT, MA, TS, and UT per 5-day trading period.

In table 3, annualized risk–return measures for all sectors are depicted. We observe annualized mean returns after transaction costs ranging from 31.90 percent for CS to 278.61 percent for FI — the naive buy–and–hold strategy is clearly outperformed (1.76 percent). This picture barely changes considering the Sharpe ratio, i.e., the excess return per unit of deviation, since all sectors attain approximately the standard deviation of the market. Most sectors generate downside deviations below the market — a favorable effect for investors because volatility is largely driven by upside deviations. The lower partial moment risk results in Sortino ratios above 2.94, compared to 0.12 for the general market. The results based on committed and employed capital show only marginal deviations — this fact is not astonishing since the top pairs trade in almost all cases (table 2). Summarizing, we may carefully conjecture that our optimal pairs trading strategy based on Lévy-driven OU processes outlined in section 4 is meaningful.

	CS	CD	EN	FI	HC	IN	IT	MA	TS	UT	MKT
Mean return	0.3190	0.7176	1.6729	2.7861	0.7647	0.9931	1.6247	1.2712	0.7293	1.1717	0.0176
Mean excess return	0.2928	0.6835	1.6200	2.7112	0.7297	0.9536	1.5727	1.2263	0.6950	1.1287	-0.0027
Standard deviation	0.1745	0.2741	0.1990	0.2627	0.2448	0.2146	0.3190	0.2147	0.2916	0.1908	0.2005
Downside deviation	0.1086	0.1766	0.1073	0.1272	0.1671	0.1271	0.1883	0.1201	0.1846	0.1137	0.1441
Sharpe ratio	1.6772	2.4939	8.1403	10.3186	2.9807	4.4444	4.9297	5.7121	2.3835	5.9154	-0.0136
Sortino ratio	2.9365	4.0629	15.5841	21.9019	4.5778	7.8124	8.6290	10.5877	3.9507	10.3048	0.1218
Employed capital											
Mean return	0.3190	0.7176	1.6729	2.7861	0.7647	0.9931	1.6247	1.2712	0.7287	1.1718	0.0176
Sharpe ratio	1.6772	2.4939	8.1403	10.3186	2.9807	4.4444	4.9297	5.7121	2.3802	5.9149	-0.0136

Table 3: Annualized risk–return measures after transaction costs for the top 10 pairs and the different sectors, compared to a S&P long-only benchmark (MKT) from March 1998 until December 2015.

Given the remarkable returns of our strategies, we compare them with 200 bootstraps of random trading to check on robustness. Similar to Gatev et al. (2006), we combine the original entry and exit signals of the top pairs with two randomly chosen securities of the S&P 500 at that time. As expected, the average daily returns before transaction costs of the random trading are close to zero at -0.01 percent per day — a value well in line with Gatev et al. (2006). In contrast, our strategies produce daily returns between 0.12 and 0.54 percent across all sectors even after transaction costs (table 1) — results are far superior to the returns of random bootstrap trading. Thus, our strategy identifies temporal variations and exploits market inefficiencies.

5.2. Strategy performance across the sectors

Motivated by subsection 5.1 and the varying performance across the branches of trade, we directly compare the results from the different sectors considering model parameters and number of stocks. In this way, we aim to identify driving sources of positive returns across all sectors.

Our trading strategy outlined in subsection 4.2 relies on the mean-reversion of spread processes (Hameed and Mian 2015, Leung and Li 2015a, Lubnaua and Todorova 2015), combined with a high degree of variation creating trading opportunities (Zeng and Lee 2014, Liu et al. 2017). Strong exposure to mean-reversion provides a high process' predictability and thus, spreads generate profits from trading (Leung and Li 2015b, Stübinger and Endres 2017). Therefore, we rank the top pairs by mean-reversion speed θ in ascending order and record the ranking r_θ. Besides fast convergence to the equilibrium level, high fluctuation enables a high trading frequency. Volatile periods of a process are identified by high volatility σ, while sudden, large movements in the spread process are created by a high jump intensity λ. Thus, pairs are ranked by σ and λ, both in ascending order, and rankings r_σ and r_λ are recorded. All three rankings constitute $r'_m = r_\theta + r_\sigma + r_\lambda$, which is again ranked ascending to create the joint model parameter ranking r_m. We identify r_m as one potential source that drives returns in a positive way.

Besides the influence of model parameters on our strategies' performance, we analyze the effect of different sector sizes. The average sector size ranges approximately between 10 stocks for TS and 80 stocks for FI. Since with increasing number of stocks the number of potential pairs increases exponentially, greater sectors offer wider selection possibilities. This motivates us to create the second ranking r_s, chosen ascending with an increasing number of average included stocks.

Figure 3 presents the distribution of the branches of trade with respect to r_m and r_s in a three-dimensional bubble chart. The bubble sizes represent the magnitudes of annualized mean returns for the different sectors, compared to the general market (MKT) (see table 3). First of all, the outstanding performance across all risk–return measures of FI is strongly driven by both a large sector size and a high parameter rank, which vindicates our trading strategy. EN and IT with high annualized returns of 167.29 and 162.47 percent also range in the right chart area with pleasant parameter ranking. As expected, the sector CS, which possesses the worst rank concerning model parameters, achieves the smallest annualized return of 31.90 percent. Surprisingly, the smallest sector TS exhibits a very high rank r_m. Concluding, we find that the dissimilarity between the sector returns is mostly driven by the corresponding mean-reversion speed, volatility, jump intensity, and sector size.

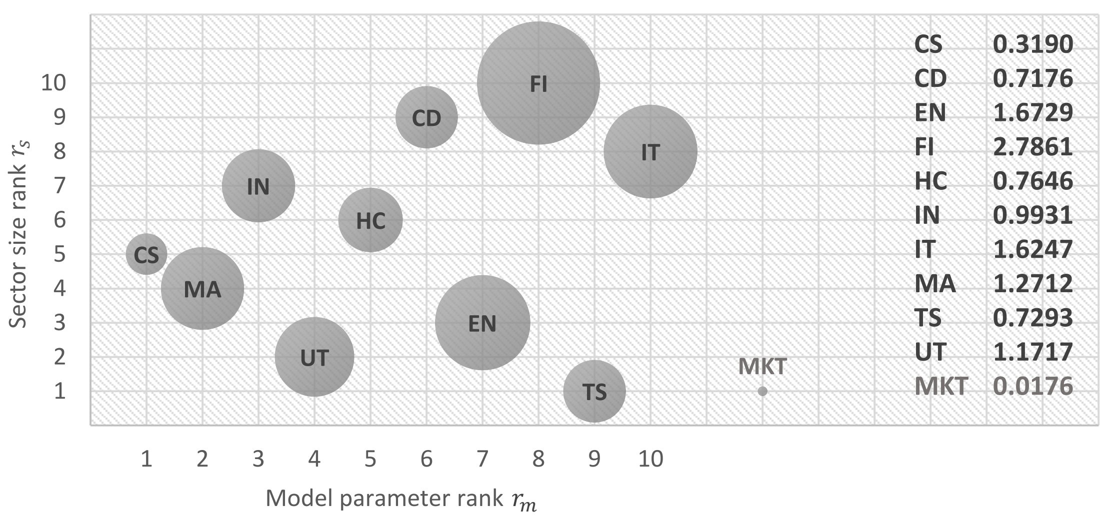

Figure 3: Performance results after transaction costs for the period from 1998 until 2015 and the different sectors in three-dimensional bubble chart. The first two dimensions visualized as coordinates are model parameter rank r_m and sector size rank r_s. The magnitudes of annualized mean returns are represented by bubble size and depicted on the right side.

5.3. Profitability in recent years

Do and Faff (2010), Bowen and Hutchinson (2016), and Stübinger and Bredthauer (2017) identify declining performance results for their pairs trading strategies in recent years. Motivated by the literature, we analyze the profitability of the strategies in the recent past.

Figure 4 presents the annualized mean return for the top 10 pairs from January 2007 to December 2015, compared to the general market (MKT). This sub-period describes the time of deterioration and moderation in stock markets caused by increasing globalization, proceeding deregulation, and faster correction of market anomalies. As expected, all 10 sectors do not achieve the average returns of the whole data sample (table 3 and figure 3), but still produce positive results after transaction costs in recent years. While most sectors generate returns between 40 and 60 percent p.a., FI outperforms clearly with an annualized return of 173.93 percent. This fact is not surprising since FI possesses both the largest sector size as well as the highest model parameters. We conclude that our optimal pairs trading strategy based on Lévy-driven OU processes generates pleasant returns even in recent past.

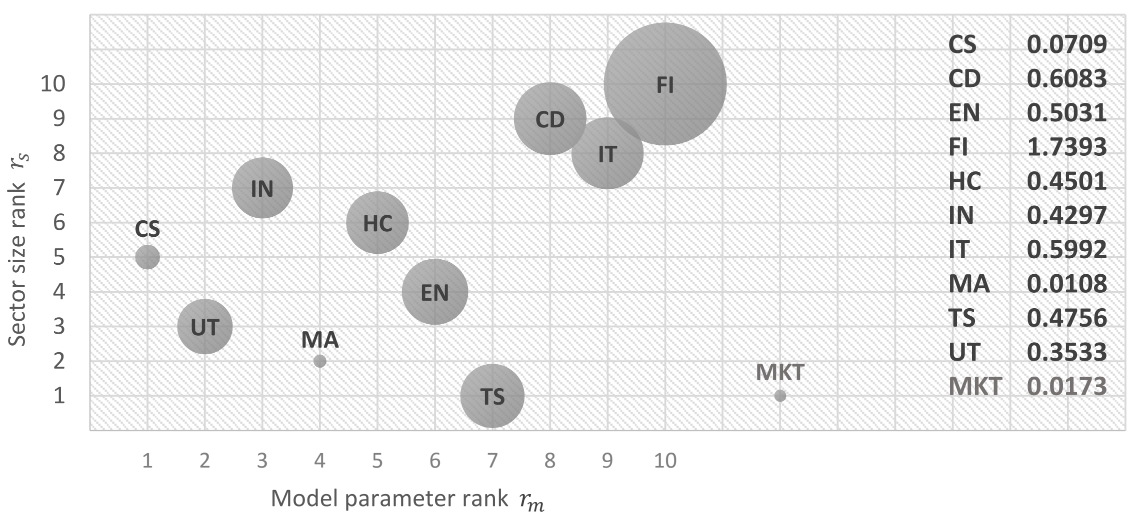

Figure 4: Performance results after transaction costs for the period from 2007 until 2015 and the different sectors in three-dimensional bubble chart. The first two dimensions visualized as coordinates are model parameter rank r_m and sector size rank r_s. The magnitudes of annualized mean returns are represented by bubble size and depicted on the right side.

5.4. Jump analysis

Johannes et al. (1999), Eraker (2004), and Kou et al. (2017) report significant changes in jump rates in equity index returns over time. Motivated by the literature, we investigate the extent of jump activity in our data.

In subsection 3.1, we specify a very flexible model to explain spread dynamics. The Lévy process is defined by

$$L_t = \sigma W_t + \sum_{i=1}^{N_t} \xi_i, \tag{10}$$

with standard Brownian motion $\{W_t\}_{t\geq 0}$, volatility $\sigma \in \mathbb{R}^+$, Poisson process $\{N_t\}_{t\geq 0}$ with rate λ, and jump sizes $\{\xi_1, \xi_2, ...\}$. As a special case for Poisson intensity $\lambda = 0$, the model covers the Brownian motion-driven OU process without any jumps. As such, our model is feasible for processes both with and without jumps. In line with this, the estimation procedure of subsection 4.1 only detects jumps if variations cannot be explained by simple Gaussian shocks.

Figure 5 shows the number of detected jumps for the different sectors over time. In a 5-day trading period, there are 1955 data points, thus 1954 time intervals and potential jumps. First of all, we observe a substantial number of detected jumps across all sectors and over the whole sample period — including jumps pays off. We confirm the finding of Eraker (2004) and observe many jumps during the first years of our data set. Especially, IT exhibits a comparatively high jump intensity during this phase, mainly affected by the dot–com bubble. The jump rates increase strongly in times of the global financial crisis from 2007 to 2009 (see Kou et al. 2017). As expected, FI contributes significantly to the high jump activity, with clearly lower jump frequencies before and after the crisis. In present days, the intensity of jumps decreases across all sectors and a major part of these jumps is caused by EN — this seems reasonable taking into account the sharp drop of oil prices beginning in the mid of 2014 and lasting until end of 2015. We may conclude that in recent years, when jump rates are at much lower levels, trading frequencies decrease and taking advantage of temporary mispricings becomes harder. Summarizing and in line with Kou et al. (2017), we observe a declining trend of jump numbers across all sectors with strong outliers in times of high market turmoils.

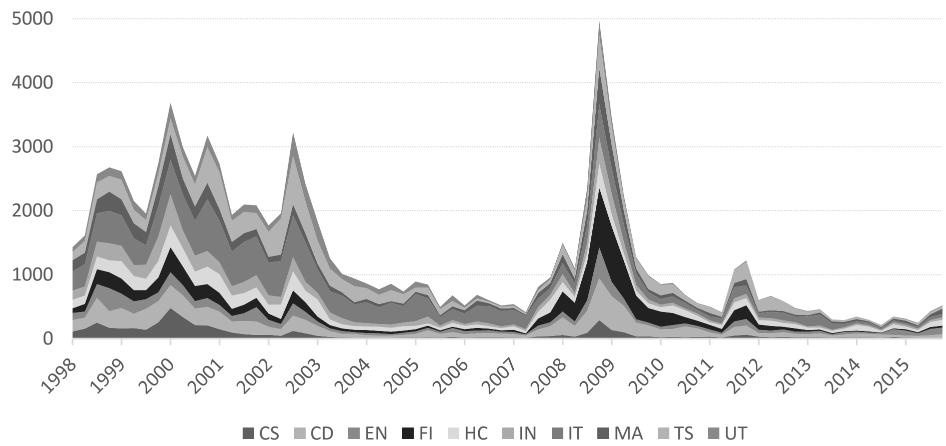

Figure 5: Average number of jumps during 5-day trading periods for the different sectors over time.

5.5. Assessment of the developed strategy

To demonstrate the additional benefit of our optimal trading strategy for Lévy-driven OU models (OLM), we compare it with well established quantitative strategies in the field of pairs trading based on stochastic differential equations. Following Bertram (2010a) and Cummins and Bucca (2012), we implement an optimal trading strategy for Brownian motion-driven OU models (OBM). Specifically, an analytic solution for the case of maximizing the expected return is derived. The second benchmark is given by Zeng and Lee (2014) and represents a classic asymmetric strategy based on Brownian motion-driven OU models (CBM). The strategy takes positions at two-standard deviations and clears positions when the spread reverts back to the mean (Bollinger 1992, 2001, Avellaneda and Lee 2010, Clegg and Krauss 2016, Stübinger et al. 2016). In stark contrast to the optimal trading rules of OLM and OBM, the entry and exit signals in CBM are model-free. Identical to OLM, we select pairs based on highest mean-reversion speed and highest variance in both benchmarks (see subsection 4.1). Following Kanamura et al. (2010), Cummins and Bucca (2012), and Liu et al. (2017), our comparison study focuses on the energy sector and depicts the performance and risk profile for OLM, OBM, and CBM in light of strategy performance over time (subsection 5.5.1) and systematic sources of risk (subsection 5.5.2).

5.5.1. Performance over time

Figure 6 displays the development of an investment of 1 USD after transaction costs for OLM, OBM, and CBM over four sub-periods, compared to the cumulative returns of the American stock market index S&P 500 (MKT) and the crude oil price of West Texas Intermediate (WTI), the most often referenced grade in oil pricing.

The first sub-period ranges from 1998 to 2003 and describes the growth and collapse of the dot–com bubble, the September 11 attacks and the start of the Iraq war. Simultaneously, there is neither an invention of methods or algorithms used in this paper nor the awareness level of high-frequency trading. As such, it is not surprising that the pairs trading strategies show strong and consistent outperformance. Optimal thresholds produce superior results compared to the model-free 2-σ rule, leading to annualized returns of 584.55 and 481.01 percent for OLM and OBM, and 156.81 percent for CBM.

The second sub-period ranges from 2004 to 2006 and characterizes the time of moderation and rising costs of oil in consequence of several events, e.g., the destruction caused by hurricane Katrina and the reducing strength of the U.S. dollar. We observe increasing developments for all time series with OLM showing particularly high performance. Specifically, our strategy outlined in section 4 generates 143.16 percent p.a..

The third sub-period ranges from 2007 to 2012 and corresponds with the global financial crisis and its aftermath. The S&P 500 index as well as the crude oil price are strongly affected by this crash — the S&P 500 loses approximately 50 percent of its value and the oil price decreases fourfold in this time. In stark contrast, OLM, OBM, and CBM generate strongly positive returns, mostly driven by the long–short portfolios we are constructing.

The forth sub-period ranges from 2013 to 2015 and specifies the period of comebacks and 2010s oil glut. The pairs trading benchmarks OBM and CBM show declining trends in comparison to the general market caused by increasing public availability of these methods. Most notable, OLM exhibits a clear growth in value up to 2.5 during the years 2013 to 2015, after transaction costs.

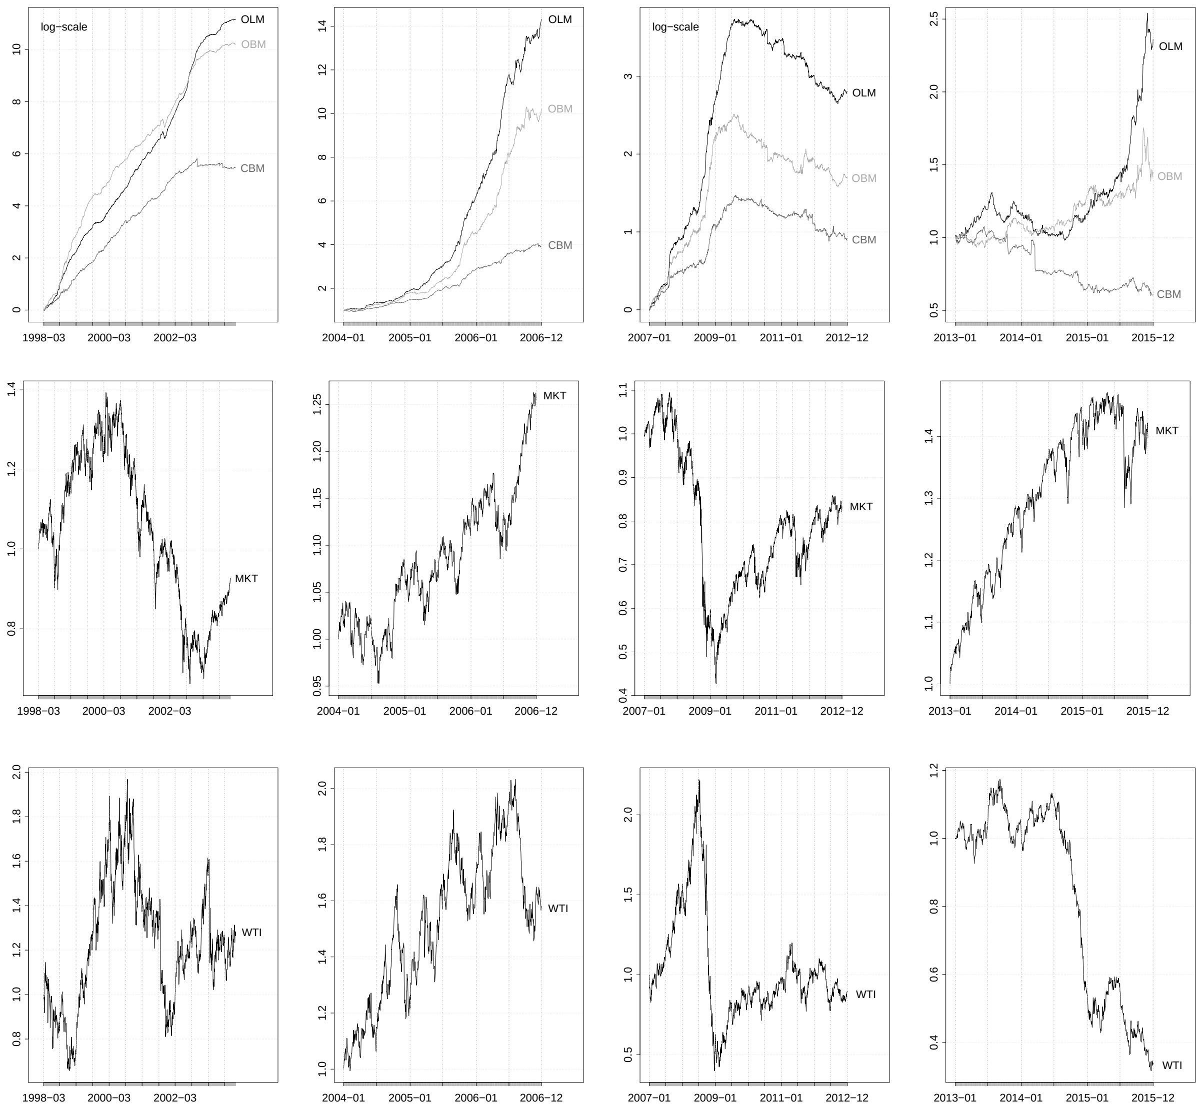

Figure 6: Development of an investment of 1 USD after transaction costs for the top 10 pairs of OLM, OBM, and CBM in the first row, compared to the S&P 500 index (MKT) in the second row and the crude oil price (WTI) in the third row. The time period from 1998 until 2015 is split into four sub-periods (1998–2003, 2004–2006, 2007–2012, 2013–2015). In the top left corner of each plot log-scales are indicated where applied.

5.5.2. Common risk factors

Finally, table 4 investigates the exposure of OLM, OBM, and CBM to common systematic sources of risk. We follow Krauss and Stübinger (2017) and apply (i) the Fama–French 3-factor model (FF3) introduced by Fama and French (1996), (ii) the Fama–French 3+2-factor model (FF3+2) outlined in Gatev et al. (2006), and (iii) the Fama–French 5-factor model (FF5) discussed by Fama and French (2015). FF3 measures systematic risk exposure to general market, small minus big capitalization stocks (SMB), and high minus low book–to–market stocks (HML). FF3+2 augments the first model by a momentum factor and a short-term reversal factor. FF5 appends two additional factors to the FF3, namely portfolios of stocks with robust minus weak profitability (RMW) and conservative minus aggressive investment behavior (CMA). All data related to these models are downloaded from Kenneth R. French's website[2].

Irrespective of the model employed, we observe that returns depict statistically and economically significant daily alphas of 0.39 percent for OLM, 0.33 percent for OBM, and 0.16 percent for CBM. Across all strategies, loading on the reversal factor is well expressed and highly significant, confirming the effect of mean-reversion that our strategy relies on. Additionally, we observe small and insignificant loadings on the remaining factors for OLM. In stark contrast, OBM and CBM capture systematic risk by possessing significant and negative loadings on the market and SMB. Overall, OLM generates statistically significant and economically considerable daily intercept alphas of 0.39 percent, does not load on any common sources of systematic risk, and outperforms the benchmark approaches — complexity pays off.

[2]We thank Kenneth R. French for all relevant data for these models on his website.

	OLM			OBM			CBM		
	FF3	FF3+2	FF5	FF3	FF3+2	FF5	FF3	FF3+2	FF5
(Intercept)	0.0039***	0.0039***	0.0039***	0.0033***	0.0032***	0.0033***	0.0016***	0.0016***	0.0016***
	(0.0002)	(0.0002)	(0.0002)	(0.0002)	(0.0002)	(0.0002)	(0.0002)	(0.0002)	(0.0002)
Market	0.0026	−0.0242	−0.0027	−0.0125	−0.0457**	−0.0188	0.0071	−0.0120	0.0096
	(0.0148)	(0.0164)	(0.0171)	(0.0158)	(0.0175)	(0.0183)	(0.0130)	(0.0144)	(0.0151)
SMB	−0.0353	−0.0282		−0.0674*	−0.0619		−0.0711**	−0.0731**	
	(0.0298)	(0.0299)		(0.0319)	(0.0320)		(0.0262)	(0.0262)	
HML	−0.0058	−0.0148		−0.0453	−0.0413		−0.0358	−0.0101	
	(0.0280)	(0.0300)		(0.0299)	(0.0321)		(0.0246)	(0.0263)	
Momentum		−0.0385			−0.0229			0.0252	
		(0.0209)			(0.0223)			(0.0183)	
Reversal		0.0668**			0.1032***			0.0913***	
		(0.0210)			(0.0225)			(0.0184)	
SMB5			−0.0400			−0.0733*			−0.0521
			(0.0322)			(0.0344)			(0.0283)
HML5			0.0094			−0.0149			−0.0134
			(0.0318)			(0.0340)			(0.0279)
RMW5			−0.0172			−0.0020			0.0439
			(0.0415)			(0.0444)			(0.0365)
CMA5			−0.0185			−0.0506			−0.0415
			(0.0509)			(0.0544)			(0.0447)
R^2	0.0003	0.0035	0.0004	0.0015	0.0066	0.0019	0.0019	0.0076	0.002
Adj. R^2	-0.0004	0.0024	-0.0007	0.0008	0.0055	0.0008	0.0012	0.0064	0.0009
Num. obs.	4484	4484	4484	4484	4484	4484	4484	4484	4484
RMSE	0.0125	0.0125	0.0125	0.0134	0.0134	0.0134	0.011	0.011	0.011

$^{***}p < 0.001$, $^{**}p < 0.01$, $^{*}p < 0.05$

Table 4: Exposure to systematic sources of risk for the daily returns of the top 10 pairs of OLM, OBM, and CBM after transaction costs from March 1998 until December 2015. Standard errors are depicted in parentheses.

6. Conclusion

In this paper, we introduce an optimal pairs trading framework based on Lévy-driven Ornstein–Uhlenbeck (OU) processes and apply it to high-frequency data of the S&P 500 constituents from January 1998 to December 2015. In respect thereof, we make four main contributions to the existing literature.

The first contribution bears on the developed optimal trading strategy for Lévy-driven OU processes. Concerning this matter, we present the first manuscript deriving an explicit representation of the expected return per unit time which serves as target function in our jump-based strategy. This takes place without the necessity of Monte Carlo approaches or numerical methods. To be more specific, the expected first-passage time is formulated by means of a mathematical expression, allowing for direct optimization of the objective function. In this way, we obtain the desired trade thresholds. By generalizing the classic Gaussian-driven OU model, we are in a position to capture additionally jumps and fat tails — both depict typical characteristics of financial data.

The second contribution refers to our high-frequency back-testing study of the S&P 500 constituents from 1998 to 2015. We observe statistically and economically significant returns after transaction costs for the top 10 pairs across all sectors demonstrating the efficiency of our strategy. Specifically, annualized returns range between 31.90 percent for the sector "Consumer Staples" and 278.61 percent for the sector "Financials", compared to 1.76 percent for a naive buy–and–hold strategy of the S&P 500. We find that the profits of all branches of trade are not being arbitraged away in the recent past.

The third contribution focuses on the varying performance results between the branches of trade. Our results reveal that there are common effects influencing returns across all economic sectors. We find that the dissimilarity is mostly driven by the respective mean-reversion speed, volatility, jump intensity, and sector size.

The forth contribution relies on the value-add of our optimal trading strategy for Lévy-driven OU models in comparison to well established Brownian motion-driven approaches in this area of research. We find that the jump-based pairs trading framework outperforms classic strategies in light of performance over time and the exposure to systematic sources of risk.

For future research in this field, time-changed Lévy processes may be explored to capture non-normal return innovations, stochastic volatility, and leverage effects. However, with increasing model complexity, the first-passage time problem becomes more challenging. Second, simultaneous trading of more than two stocks which co-move in some pattern could be executed. Therefore, a multivariate model that accounts for common interactions has to be estimated. Third, interdependencies and dissimilarities among the sector returns should be elaborated and explained by additional influencing factors.

Bibliography

Aït-Sahalia, Y., Jacod, J., 2014. High-frequency financial econometrics. Princeton University Press, New Jersey, USA.

Avellaneda, M., Lee, J.-H., 2010. Statistical arbitrage in the US equities market. Quantitative Finance 10 (7), 761–782.

Bacon, C. R., 2008. Practical portfolio performance: Measurement and attribution, 2nd Edition. John Wiley & Sons, Chichester, England.

Bai, Y., Wu, L., 2017. Analytic value function for optimal regime-switching pairs trading rules. Quantitative Finance, forthcoming.

Barlow, M. T., 2002. A diffusion model for electricity prices. Mathematical Finance 12 (4), 287–298.

Baviera, R., Baldi, T. S., 2017. Stop-loss and leverage in optimal statistical arbitrage with an application to energy market. Working paper, Politecnico di Milano.

Bayraktar, E., Xing, H., 2011. Pricing Asian options for jump diffusion. Mathematical Finance 21 (1), 117–143.

Bertram, W. K., 2009. Optimal trading strategies for Itô diffusion processes. Physica A: Statistical Mechanics and its Applications 388 (14), 2865–2873.

Bertram, W. K., 2010a. Analytic solutions for optimal statistical arbitrage trading. Physica A: Statistical Mechanics and its Applications 389 (11), 2234–2243.

Bertram, W. K., 2010b. A non-stationary model for statistical arbitrage trading. Working paper, ITG Australia Limited.

Bogomolov, T., 2013. Pairs trading based on statistical variability of the spread process. Quantitative Finance 13 (9), 1411–1430.

Bollinger, J., 1992. Using Bollinger bands. Stocks & Commodities 10 (2), 47–51.

Bollinger, J., 2001. Bollinger on Bollinger bands. McGraw-Hill, New York, USA.

Borovkov, K., Novikov, A., 2008. On exit times of Lévy-driven Ornstein–Uhlenbeck processes. Statistics & Probability Letters 78 (12), 1517–1525.

Bowen, D. A., Hutchinson, M. C., 2016. Pairs trading in the UK equity market: Risk and return. The European Journal of Finance 22 (14), 1363–1387.

Cai, N., Kou, S. G., 2011. Option pricing under a mixed-exponential jump diffusion model. Management Science 57 (11), 2067–2081.

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

Carr, P., Wu, L., 2004. Time-changed Lévy processes and option pricing. Journal of Financial Economics 71 (1), 113–141.

Cartea, Á., Figueroa, M. G., 2005. Pricing in electricity markets: A mean reverting jump diffusion model with seasonality. Applied Mathematical Finance 12 (4), 313–335.

Clegg, M., Krauss, C., 2016. Pairs trading with partial cointegration. FAU Discussion Papers in Economics (5), University of Erlangen-Nürnberg.

Cont, R., 2001. Empirical properties of asset returns: Stylized facts and statistical issues. Quantitative Finance 1 (2), 223–236.

Cont, R., Mancini, C., 2011. Nonparametric tests for pathwise properties of semimartingales. Bernoulli 17 (2), 781–813.

Cont, R., Tankov, P., 2003. Financial modelling with jump processes. Vol. 2. CRC Press, London, England.

Cummins, M., Bucca, A., 2012. Quantitative spread trading on crude oil and refined products markets. Quantitative Finance 12 (12), 1857–1875.

Do, B., Faff, R., 2010. Does simple pairs trading still work? Financial Analysts Journal 66 (4), 83–95.

Do, B., Faff, R., 2012. Are pairs trading profits robust to trading costs? Journal of Financial Research 35 (2), 261–287.

Ekström, E., Lindberg, C., Tysk, J., 2011. Optimal liquidation of a pairs trade. In: Di Nunno, G., Øksendal, B. (Eds.), Advanced mathematical methods for finance. Springer, Berlin, Heidelberg, pp. 247–255.

Elliott, R. J., van der Hoek, J., Malcolm, W. P., 2005. Pairs trading. Quantitative Finance 5 (3), 271–276.

Eraker, B., 2004. Do stock prices and volatility jump? Reconciling evidence from spot and option prices. The Journal of Finance 59 (3), 1367–1403.

Fama, E. F., French, K. R., 1996. Multifactor explanations of asset pricing anomalies. The Journal of Finance 51 (1), 55–84.

Fama, E. F., French, K. R., 2015. A five-factor asset pricing model. Journal of Financial Economics 116 (1), 1–22.

Fox, P., 1997. The port mathematical subroutine library, version 3. AT&T Bell Laboratories, Murray Hill.

Gatev, E., Goetzmann, W. N., Rouwenhorst, K. G., 1999. Pairs trading: Performance of a relative value arbitrage rule. Working paper, Yale School of Management's International Center for Finance.

Gatev, E., Goetzmann, W. N., Rouwenhorst, K. G., 2006. Pairs trading: Performance of a relative-value arbitrage rule. Review of Financial Studies 19 (3), 797–827.

Gay, D. M., 1990. Usage summary for selected optimization routines. Computing science technical report 153, 1–21.

Gerber, H. U., Shiu, E. S. W., Yang, H., 2013. Valuing equity-linked death benefits in jump diffusion models. Insurance: Mathematics and Economics 53 (3), 615–623.

Göncü, A., Akyıldırım, E., 2016a. Statistical arbitrage with pairs trading. International Review of Finance 16 (2), 307–319.

Göncü, A., Akyıldırım, E., 2016b. A stochastic model for commodity pairs trading. Quantitative Finance 16 (12), 1843–1857.

Govender, K., 2011. Statistical arbitrage in South African financial markets. Working paper, University of Cape Town.

Gregory, I., Ewald, C.-O., Knox, P., 2011. Analytical pairs trading under different assumptions on the spread and ratio dynamics. Working paper, University of Sydney.

Hameed, A., Mian, G. M., 2015. Industries and stock return reversals. Journal of Financial and Quantitative Analysis 50 (1-2), 89–117.

Huang, J.-Z., Huang, M., 2012. How much of the corporate-treasury yield spread is due to credit risk? The Review of Asset Pricing Studies 2 (2), 153–202.

Huck, N., Afawubo, K., 2015. Pairs trading and selection methods: Is cointegration superior? Applied Economics 47 (6), 599–613.

Jegadeesh, N., Titman, S., 1993. Returns to buying winners and selling losers: Implications for stock market efficiency. The Journal of Finance 48 (1), 65–91.

Jing, B.-Y., Kong, X.-B., Liu, Z., 2012. Modeling high-frequency financial data by pure jump processes. Annals of Statistics, 759–784.

Johannes, M., Kumar, R., Polson, N. G., 1999. State dependent jump models: How do US equity indices jump? Woking Paper, University of Chicago.

Jondeau, E., Lahaye, J., Rockinger, M., 2015. Estimating the price impact of trades in a high-frequency microstructure model with jumps. Journal of Banking & Finance 61 (2), 205–224.

Kanamura, T., Rachev, S. T., Fabozzi, F. J., 2010. A profit model for spread trading with an application to energy futures. The Journal of Trading 5 (1), 48–62.

Knoll, J., Stübinger, J., Grottke, M., 2017. Exploiting social media with higher-order factorization machines: Statistical arbitrage on high-frequency data of the S&P 500. FAU Discussion Papers in Economics (13), University of Erlangen-Nürnberg.

Kou, S., Yu, C., Zhong, H., 2017. Jumps in equity index returns before and during the recent financial crisis: A Bayesian analysis. Management Science 63 (4), 988–1010.

Kou, S. G., 2002. A jump-diffusion model for option pricing. Management Science 48 (8), 1086–1101.

Kou, S. G., Wang, H., 2003. First passage times of a jump diffusion process. Advances in Applied Probability 35 (2), 504–531.

Krauss, C., Stübinger, J., 2017. Non-linear dependence modelling with bivariate copulas: Statistical arbitrage pairs trading on the S&P 100. Applied Economics 49 (52), 5352–5369.

Larsson, S., Lindberg, C., Warfheimer, M., 2013. Optimal closing of a pair trade with a model containing jumps. Applications of Mathematics 58 (3), 249–268.

Leung, T., Li, X., 2015a. Optimal mean reversion trading: Mathematical analysis and practical applications. World Scientific, New Jersey, USA.

Leung, T., Li, X., 2015b. Optimal mean reversion trading with transaction costs and stop-loss exit. International Journal of Theoretical and Applied Finance 18 (3), 1550020.

Li, H., Wells, M. T., Yu, C. L., 2008. A Bayesian analysis of return dynamics with Lévy jumps. Review of Financial Studies 21 (5), 2345–2378.

Li, X., 2015. Optimal multiple stopping approach to mean reversion trading. PhD thesis, Columbia University.

Liu, B., Chang, L.-B., Geman, H., 2017. Intraday pairs trading strategies on high frequency data: The case of oil companies. Quantitative Finance 17 (1), 87–100.

Lubnaua, T., Todorova, N., 2015. Trading on mean-reversion in energy futures markets. Energy Economics 51, 312–319.

Mai, D.-M. H., 2012. Drift estimation for jump diffusions: Time-continuous and high-frequency observations. Ph.D. thesis, Humboldt-Universität zu Berlin.

Mai, H., 2014. Efficient maximum likelihood estimation for Lévy-driven Ornstein–Uhlenbeck processes. Bernoulli 20 (2), 919–957.

Mancini, C., 2009. Non–parametric threshold estimation for models with stochastic diffusion coefficient and jumps. Scandinavian Journal of Statistics 36 (2), 270–296.

Meyer-Brandis, T., Tankov, P., 2008. Multi-factor jump-diffusion models of electricity prices. International Journal of Theoretical and Applied Finance 11 (5), 503–528.

Miao, G. J., 2014. High frequency and dynamic pairs trading based on statistical arbitrage using a two-stage correlation and cointegration approach. International Journal of Economics and Finance 6 (3), 96–110.

Mina, J., Xiao, J. Y., 2001. Return to RiskMetrics: The evolution of a standard. RiskMetrics Group.

Peterson, B. G., Carl, P., 2014. PerformanceAnalytics: Econometric tools for performance and risk analysis. R package.

QuantQuote, 2016. QuantQuote market data and software. https://www.quantquote.com/.

R Core Team, 2017. R package stats: A language and environment for statistical computing. R package.

Ramezani, C. A., Zeng, Y., 1998. Maximum likelihood estimation of asymmetric jump-diffusion processes: Application to security prices. Working paper, California Polytechnic State University.

Ramezani, C. A., Zeng, Y., 2007. Maximum likelihood estimation of the double exponential jump-diffusion process. Annals of Finance 3 (4), 487–507.

Ryan, J. A., Ulrich, J. M., 2014. xts: eXtensible time series. R package.

Sepp, A., 2003. Fourier transform for option pricing under affine jump-diffusions: An overview. Working paper, University of Tartu.

S&P 500 Dow Jones Indices, 2015. Equity S&P 500. http://us.spindices.com/indices/equity/sp-500.

Stübinger, J., Bredthauer, J., 2017. Statistical arbitrage pairs trading with high-frequency data. International Journal of Economics and Financial Issues 7 (4), 650–662.

Stübinger, J., Endres, S., 2017. Pairs trading with a mean-reverting jump-diffusion model on high-frequency data. FAU Discussion Papers in Economics (10), University of Erlangen-Nürnberg.

Stübinger, J., Mangold, B., Krauß, C., 2016. Statistical arbitrage with vine copulas. FAU Discussion Papers in Economics (11), University of Erlangen-Nürnberg.

Su, X., Bai, M., 2017. First-passage time model driven by Lévy process for pricing CoCos. Mathematical Problems in Engineering (2017), 1–13.

Suzuki, K., 2017. Optimal pair-trading strategy over long/short/square positions—empirical study. Quantitative Finance, forthcoming.

Tsay, R. S., 2005. Analysis of financial time series. John Wiley & Sons, New Jersey, USA.

Ulrich, J., 2016. TTR: Technical trading rules. R package.

Vidyamurthy, G., 2004. Pairs trading: Quantitative methods and analysis. John Wiley & Sons, Hoboken, USA.

Zeng, Z., Lee, C.-G., 2014. Pairs trading: Optimal thresholds and profitability. Quantitative Finance 14 (11), 1881–1893.