Quantitative Finance, 2016
http://dx.doi.org/10.1080/14697688.2016.1231928

Routledge
Taylor & Francis Group

Model-based pairs trading in the bitcoin markets

P. S. LINTILHAC† and A. TOURIN*‡

†Department of Mathematics, New York University Tandon School of Engineering, Six Metrotech Center, Brooklyn, NY 11201, USA

‡Department of Finance and Risk Engineering, New York University Tandon School of Engineering, Six Metrotech Center, Brooklyn, NY 11201, USA

(*Received 27 February 2016; accepted 25 August 2016; published online 4 November 2016*)

We propose an optimal dynamic pairs trading strategy model for a portfolio of cointegrated assets. Using stochastic control techniques, we compute analytically the optimal portfolio weights and relate our result to several other strategies commonly used by practitioners, including the static double-threshold strategy. Finally, we apply our model to a bitcoin portfolio and conduct an out-of-sample test with historical data from three exchanges, with two cointegrating relations.

Keywords: Stochastic control; Pairs trading; Merton problem; Cointegration; Statistical arbitrage

JEL Classification: C61, G11

1. Introduction

This article generalizes the dynamic pairs trading model in Tourin and Yan (2013) to a portfolio of cointegrated assets of arbitrary size and applies it to bitcoin markets. The model, which characterizes the optimal asset holdings dynamically, can be written in vector form and solved analytically using stochastic control techniques.

Our approach combines the model of cointegrated assets by Duan and Pliska (2004) with the Merton portfolio selection problem (Merton 1971). In contrast to Tourin and Yan (2013), we include a linear trend in the cointegrating relations and make the dynamics fully symmetric. The investor's objective is to maximize the expected utility derived from terminal wealth. In this work, we choose an exponential utility function for simplicity of calculation. Using classical stochastic control techniques, we compute the optimal portfolio weights in closed form and validate this result by proving a verification Theorem which provides an upper bound on the time horizon in terms of the parameters in the model beyond which the solution may blow up. We also briefly explore several other alternate strategies, including the double-threshold strategy which is commonly used by practitioners. Finally, we illustrate the applicability of our method by conducting both an in-sample and an out-of-sample test with market data from three bitcoin exchanges.

The application of stochastic control to pairs trading originated in the work of Mudchanatongsuk *et al.* (2008). More recently, two articles formulated a model characterizing the optimal entry and exit points of a pairs trading strategy. The first one by Leung and Li (2015) also incorporates transaction costs. The second one by Lei and Xu (2015) determines multiple entry and exit-points during a trading period and includes a performance study carried out on dual-listed Chinese stocks. Besides, Ngo and Pham (2016) frame the pairs trading problem as a regime switching model between three regimes: flat positions, one long position on one asset and a short position on the other, and vice versa. Portfolios in several dimensions have also been considered: among others, a static mean-variance approach based on the cointegration model by Duan and Pliska (2004), has been developed by Chiu and Wong (2011) while Avellaneda and Lee (2010) provided an empirical study for beta-neutral pairs trading portfolio strategies in the US equities market. Besides, we refer to the books by Fleming and Soner (1993) and Pham (2009) for an introduction to stochastic control and its applications in Finance. Finally, during the preparation of this manuscript, it came to our attention that Cartea and Jaimungal (forthcoming) also investigated the multi-dimensional version of the model of Tourin and Yan (2013), producing results that complement ours.

In the second section, we present the optimal stochastic control problem, derive its solution and present a verification result. In section 3, we investigate several alternate formulations, including the well-known double-threshold strategy. Finally, in section 4, we show our experimental results in the bitcoin markets.

*Corresponding author. Email: atourin@nyu.edu

2. The general model for a portfolio of cointegrated and correlated assets

First of all, we present the model. Then, we compute its solution in the second subsection. In the last subsection, we provide a verification Theorem ensuring that the computed strategy is the solution of the original stochastic control problem.

2.1. Problem formulation

We consider a portfolio of n cointegrated assets and model the evolution of their log-prices as in Duan and Pliska (2004). We allow the assets to be linearly correlated and we assume the volatility matrix to be constant. Specifically, the asset prices are driven by an n-dimensional standard Brownian motion $B_t = (B_t^1, \dots, B_t^n)$. The vector of asset prices is denoted by $S_t = (S_t^1, \dots, S_t^n)^t$, and $x_t = (x_t^1, \dots, x_t^n)^t = (\log S_t^1, \dots, \log S_t^n)^t$ represents the vector of log-prices. We say that the assets are cointegrated with exactly r cointegrating relations when there is a $n \times r$ matrix β of rank r, where $0 < r < n$, and r-dimensional column vectors a, b, such that $a + bt + \beta^t x_t$ is an Ornstein–Uhlenbeck process. We denote the r-dimensional spread by $z_t = (z_t^1, \dots, z_t^r)^t$, i.e.

$$z_t = a + bt + \beta^t x_t, \tag{1}$$

where β^t denotes the transpose matrix of β. Note that the term bt can be used to offset the potential deterministic trend in $\beta^t x_t$, in order to make z_t an Ornstein–Uhlenbeck process with long-run mean 0.

Next, the log-prices follow the dynamics

$$\mathrm{d}x_t = \left(\mu - \frac{1}{2} D(\sigma\sigma^t) + \delta z_t\right) \mathrm{d}t + \sigma \mathrm{d}B_t, \tag{2}$$

where $\mu = (\mu_1, \dots \mu_n)$ is the $n \times 1$ vector of drift coefficients, $\sigma = (\sigma_{ij})$ is a $n \times n$ invertible matrix, δ is a $n \times r$ matrix of rank r, $D(\sigma\sigma^t)$ is the vector whose elements are the diagonal of $\sigma\sigma^t$.

Moreover, it follows from (1) and (2) that z_t satisfies the dynamics

$$\begin{aligned} dz_t &= \left(b + \beta^t \mu - \frac{1}{2}\beta^t D(\sigma\sigma^t) + \beta^t \delta z_t\right) \mathrm{d}t + \beta^t \sigma \mathrm{d}B_t, \\ &= -\beta^t \delta(\nu - z_t)\mathrm{d}t + \beta^t \sigma \mathrm{d}B_t, \end{aligned}$$

where $\nu = -(\beta^t\delta)^{-1}(b + \beta^t\mu - \frac{1}{2}\beta^t D(\sigma\sigma^t))$ is the long-run mean.

Furthermore, according to Duan and Pliska (2004), the system (1), (2) is cointegrated if and only if

$$||I_r + \beta^t \delta|| < 1.$$

In the above condition, I_r is the $r \times r$ identity matrix and $||\cdot||$ is the matrix norm

$$||A|| = \max_{x \neq 0} \frac{|Ax|}{|x|},$$

where $|x|$ is the Euclidian norm of the r-vector x.

Finally, we will remedy the over parametrization of (1), (2) in section 4, in conjunction with the parameters estimation procedure.

We construct a self-financing portfolio invested in the cointegrated assets and a bank account. For simplicity, we assume that the bank account's interest rate is equal to 0. The evolution of the wealth variable reads

$$\mathrm{d}W_s = \sum_{i=1}^{n} \pi_s^i \frac{\mathrm{d}S_s^i}{S_s^i},$$

where $\pi_s = (\pi_s^1, \dots, \pi_s^n)^t$ is a $n \times 1$ vector, whose component π_s^i represents the amount invested in the ith asset at time s.

Hence, the state variables (W_s, x_s) evolve according to

$$\mathrm{d}W_s = \pi_s^t \left((\mu + \delta z_s)\mathrm{d}s + \sigma \mathrm{d}B_s\right), \tag{3}$$

$$\mathrm{d}x_s = \left(\mu - \frac{1}{2} D(\sigma\sigma^t) + \delta z_s\right) \mathrm{d}s + \sigma \mathrm{d}B_s, \tag{4}$$

$$W_t = w, x_t = x. \tag{5}$$

where z_s is defined in (1). Next, we fix a finite time horizon $T > 0$.

For the sake of simplicity, we treat the case of the exponential utility function, i.e.

$$U(w) = -e^{-\gamma w},$$

where $\gamma > 0$ denotes the constant risk aversion coefficient.

A vector π of controls is said to be admissible if the elements of π are real-valued, progressively measurable, π is such that, (1), (3)–(5) define a unique solution (W_s, x_s) for every time $s \in [0, T]$, π satisfy the integrability condition

$$\mathbb{E} \int_t^T \sum_{i=1}^{n} (\pi_s^i)^2 \mathrm{d}s < +\infty, \tag{6}$$

and the family of random variables $\exp(-\gamma W_{s\wedge\tau})$, where τ is a stopping time in $[t, T]$, is uniformly integrable.

We denote the set of admissible controls at the initial time of investment t, by $\mathcal{A}_t$. Next, we define the value function $u(t, w, x)$ of the following stochastic control problem: the investor seeks an admissible strategy π that maximizes the utility he derives from his terminal wealth at time T, i.e.

$$u(t, w, x) = \sup_{\pi \in \mathcal{A}_t} \mathbb{E}[U(W_T^{t,w,x,\pi})], \tag{7}$$

where $W_T^{t,w,x,\pi}$ denotes the solution of (3) at time T, corresponding to the control π, the stock log-prices x whose dynamics are defined by (1), (4) and for the initial conditions (5).

The value function $u(t, w, x)$ satisfies the HJB equation

$$\begin{aligned} -u_t - \sup_{\pi} \Big[&\{\pi^t(\mu + \delta z)\} u_w \\ &+ \left(\mu - \frac{1}{2} D(\sigma\sigma^t) + \delta z\right) \cdot D_x u + \sigma\sigma^t \pi \cdot D_x u_w \\ &+ \frac{1}{2}\pi^t \sigma\sigma^t \pi u_{ww} + \frac{1}{2}\mathrm{tr}\left(\sigma\sigma^t D_x^2 u\right) \Big] = 0, \end{aligned} \tag{8}$$

for all $0 \leq t < T$, $(w, x) \in \mathbb{R}^{n+1}$, where $z = z(t, x) = a + bt + \beta^t x$, and the terminal condition

$$u(T, w, x) = -e^{-\gamma w}, \quad \text{for all } (w, x) \in \mathbb{R}^{n+1}, \tag{9}$$

2.2. The solution

In what follows, we solve explicitly the HJB equations (8), (9). In order to factor out the wealth variable and to reduce the number of spatial variables, we substitute the Ansatz

$$u(t,w,x) = -e^{-\gamma w} h(t,z),$$

where

$$z = a + bt + \beta^t x, \tag{10}$$

into (8), (9), and derive the PDE characterizing h:

$$\begin{aligned} -h_t + \sup_{\pi}\Big[&\{\pi^t(\mu+\delta z)\}\gamma h - \beta^t\left(\mu - \frac{1}{2}D(\sigma\sigma^t) + \delta z\right)\cdot D_z h \\ &+ \gamma\beta^t\sigma\sigma^t\pi\cdot D_z h - b\cdot D_z h - \frac{1}{2}\gamma^2\pi^t\sigma\sigma^t\pi h \\ &- \frac{1}{2}\mathrm{tr}\left(\sigma\sigma^t\beta D^2_{zz}h\beta^t\right)\Big] = 0, \end{aligned} \tag{11}$$

for all $0 \le t < T$, $z \in \mathbb{R}^r$, coupled with

$$h(T,z) = 1, \quad \text{for all } z \in \mathbb{R}^r.$$

We solve analytically the maximization problem in (11) to compute the controls in terms of h and its partial derivatives in feedback form.

$$\pi^*(t,z) = \frac{1}{\gamma}\left((\sigma\sigma^t)^{-1}(\mu+\delta z) + \frac{1}{h(t,z)}\beta D_z h(t,z)\right).$$

Substituting this expression back into (11), we obtain the PDE

$$\begin{aligned} -h_t &+ \frac{1}{2}(\mu+\delta z)^t(\sigma\sigma^t)^{-1}(\mu+\delta z)h + \frac{1}{2}\beta^t D(\sigma\sigma^t)\cdot D_z h \\ &+ \frac{1}{2h}D_z h^t\beta^t\sigma\sigma^t\beta D_z h - b\cdot D_z h \\ &- \frac{1}{2}\mathrm{tr}\left(\sigma\sigma^t\beta D^2_{zz}h\beta^t\right)\Big] = 0, \quad \text{for all } t\in[0,T), z\in\mathbb{R}^r. \end{aligned} \tag{12}$$

We get rid of the nonlinearity in (12), by making the change of unknown function

$$h(t,z) = \exp(-\phi(t,z)),$$

and deriving the linear PDE satisfied by ϕ

$$\begin{aligned} -\phi_t &- \frac{1}{2}(\mu+\delta z)^t(\sigma\sigma^t)^{-1}(\mu+\delta z) \\ &+ \frac{1}{2}\beta^t D(\sigma\sigma^t)\cdot D_z\phi - b\cdot D_z\phi \\ &- \frac{1}{2}\mathrm{tr}\left(\sigma\sigma^t\beta D^2_{zz}\phi\beta^t\right)\Big] = 0, \quad \text{for all } t\in[0,T), z\in\mathbb{R}^r, \end{aligned} \tag{13}$$

coupled with the terminal condition

$$\phi(T,z) = 0, \quad \text{for all } z\in\mathbb{R}^r.$$

Next, we use the Ansatz

$$\phi(t,z) = z^t A(t) z + B(t) z + C(t),$$

where $A(t)$ is a $r\times r$ matrix, $B(t)$ is a $1\times r$ matrix and $C(t)$ is a real number and $A(T)=0, B(T)=0, C(T)=0$. We substitute this Ansatz into (13), and easily solve the three Ordinary Differential Equations characterizing the coefficients $A(t), B(t), C(t)$. We find

$$A(t) = \frac{1}{2}\delta^t(\sigma\sigma^t)^{-1}\delta(T-t), \tag{14}$$

$$B(t) = \mu^t(\sigma\sigma^t)^{-1}\delta(T-t) + \left(b^t - \frac{1}{2}D(\sigma\sigma^t)^t\beta\right)\delta^t(\sigma\sigma^t)^{-1}\delta\frac{(T-t)^2}{2}, \tag{15}$$

$$\begin{aligned} C(t) = &\frac{1}{2}\mu^t(\sigma\sigma^t)^{-1}\mu(T-t) \\ &- \left(\mu^t(\sigma\sigma^t)^{-1}\delta\left(\frac{1}{2}\beta^t D(\sigma\sigma^t) - b\right)\right. \\ &\left.- \mathrm{tr}(\sigma\sigma^t\beta\delta^t(\sigma\sigma^t)^{-1}\delta\beta^t)\right)\frac{(T-t)^2}{2} \\ &+ \left(\frac{1}{2}D(\sigma\sigma^t)^t\beta - b^t\right)\delta^t(\sigma\sigma^t)^{-1}\delta \\ &\times\left(\frac{1}{2}\beta^t D(\sigma\sigma^t) - b\right)\frac{(T-t)^3}{6}. \end{aligned} \tag{16}$$

Finally, the controls are

$$\pi^*(t,z) = \frac{1}{\gamma}\left((\sigma\sigma^t)^{-1}(\mu+\delta z) + \beta\left(-2A(t)z - B(t)^t\right)\right), \tag{17}$$

where $A(t)$, $B(t)$ are given in (14) and (15).

2.3. *Verification result*

Our verification argument is very similar to the one in Benth and Karlsen (2005). We only need to generalize their proof to the multi-dimensional case and adapt it to the exponential utility function. We are able to derive a bound on the time horizon, beyond which the closed-form solution we computed may blow up. We present the result below, as well as the key steps in the proof.

First of all, we define the following matrices, which are continuous functions of the time variable,

$$C_0(s) = \delta^t(\sigma\sigma^t)^{-1}\delta\left[-I_r + \beta^t\delta(T-s)\right], \tag{18}$$

$$C_1(s) = \delta^t(\sigma\sigma^t)^{-1}\left[-I_n + \delta\beta^t(T-s)\right]\sigma, \tag{19}$$

where I_r and I_n are, respectively, the $r\times r$ and $n\times n$ identity matrices. Next, we introduce the $r\times r$ covariance matrix $\Omega(s)$ of the process z_s at time s:

$$\Omega(s) = \int_t^s e^{\beta^t\delta(s-u)}\beta^t\sigma\sigma^t\beta e^{\delta^t\beta(s-u)}\mathrm{d}u. \tag{20}$$

Finally, we consider the diagonal $r\times r$ matrices, $\Lambda_0(s)$, $\Lambda_1(s)$, containing, respectively, the sets eigenvalues of $\frac{1}{2}\Omega^{\frac{1}{2}}(C_0+C_0^t)\Omega^{\frac{1}{2}}(s)$ and $\Omega^{\frac{1}{2}}C_1C_1^t\Omega^{\frac{1}{2}}(s)$.

We state our result below. The conditions in the Theorem translate into an upper bound on the time step or impose restrictions on the risk preference parameter γ.

THEOREM 2.1 *Let $T > 0$ be the given time horizon, and $t\in[0,T]$. If $4\max_{s\in[t,T]}||\Lambda_0(s)|| < 1$ and $32\max_{s\in[t,T]}||\Lambda_1(s)|| < 1$, the value function of the optimal stochastic problem is given by*

$$u(t,w,x) = -e^{-\gamma w}\exp\{-z^t A(t)z - B(t)z - C(t)\}$$

where the functions A, B, C are given in (14)–(16), z is defined in (10), and the optimal control is given in feedback form in (17).

Proof. The main point is to prove the uniform integrability of $\{u(\tau, W^*_\tau, x_\tau)\}_\tau$, where W^* is the wealth process corresponding to the optimal control (17). Following Benth and Karlsen (2005), this problem can be reduced to finding a real number $\epsilon > 0$ for which the following two expectations are finite

$$\mathbb{E}\left[\exp\left\{4(1+\epsilon)\int_t^T z_s^t C_0(s) z_s \mathrm{d}s\right\}\right], \tag{21}$$

$$\mathbb{E}\left[\exp\left\{32(1+\epsilon)^2\int_t^T z_s^t C_1(s) C_1(s)^t z_s \mathrm{d}s\right\}\right] \tag{22}$$

Next, the cointegration process can be centred as both integrals can be majorized by a constant multiplied, respectively, by

$$\mathbb{E}\left[\exp\left\{4(1+\epsilon)\int_t^T z_s^{0t} C_0(s) z_s^0 \mathrm{d}s\right\}\right], \tag{23}$$

$$\mathbb{E}\left[\exp\left\{32(1+\epsilon)^2\int_t^T z_s^{0t} C_1(s) C_1(s)^t z_s^0 \mathrm{d}s\right\}\right], \tag{24}$$

where z_s^0 is the centred cointegration process. □

We can rewrite the quadratic forms, respectively, as $y_s^t \Lambda_0(s) y_s$ and $y_s^t \Lambda_1(s) y_s$ where y_s is a $r \times 1$ vector of independent standard normal variables.

Finally, we adapt Lemma 4.3 in Benth and Karlsen (2005): The first expectation is majorized by

$$C\mathbb{E}\left[\exp\left\{4(1+\epsilon)\max_{s\in[t,T]}||\Lambda_0(s)||\int_t^T y_s^t y_s \mathrm{d}s\right\}\right], \tag{25}$$

We use the fact that $y_s^t y_s = \sum_{i=1}^r y_{is}^2$ is a centred chi-squared distribution with r degrees of freedom to conclude similarly as in Benth and Karlsen (2005). The same type of argument yields the desired result for the second expectation.

3. Connection with other strategies

We observe that the optimal control in (17) has two components: the first one is very similar to the standard Merton optimal portfolio weight, except that the process z now appears in the numerator, replacing μ by $\mu + \delta z$, as the rate of return; the second one is new and specific to the mean-reversion mechanism, in the sense that it contributes to building a cumulative profit when the prices are in a mean-reverting phase.

Some of the solution's features may be unnatural to a trader. One such characteristic is the fact that the time horizon must be known a priori. One might argue that this dependence on the time horizon makes sense since, in practice, there is often some final time at which we want to maximize our wealth, such as the end of day, and certainly during testing there is always a final time. However, in practice it is often difficult to determine exactly what the final time is, and this uncertainty translates to substantial variations in our computed solution whose second component depends explicitly the time to maturity. For example, we may be operating in a market that never closes (such as FX markets or bitcoin markets, the main example of an application discussed in this paper), or we may seek a strategy that we can simply turn on and keep running as long as possible. In cases like these, we might prefer a strategy that does not depend on the time horizon.

Examples of commonly used trading strategies for cointegrated assets are the double-threshold strategy or the relative value strategy. In both of these cases, the rules for trading depend only on the parameters of the model and the asset prices, but have no explicit time-dependence. We can explain the intuition behind the time-dependence in our model by illustrating with an example. Consider a situation where there is some significant mispricing, and the time is very far from maturity. In this situation, it makes sense that we would be willing to take on a larger position (along with more risk) because there is more time for the process to potentially revert back to the mean before the final time. In fact, if the time is long enough, it is almost guaranteed. As the time gets closer to maturity however, we are less willing to take on a large position because there is a significant chance that the prices may not revert back to the mean before maturity is reached, resulting in an unnecessary loss.

The dependence of our solution on the time to maturity is inherent to the fact that the dynamic programming principle approach is backward and eliminating the need to know a priori the time horizon would require reframing our problem. A possible solution would be to use the approach proposed by Musiela and Zariphopoulou (2008) consisting in applying a forward performance criterion. This is beyond the scope of this paper and we do not address this question herein.

In what follows, we vary slightly the assumptions in our model and observe the effects. As a first step, we replace the geometric Brownian motion by the arithmetic Brownian motion for modelling the asset prices and we assume that the price themselves are cointegrated. Secondly, we seek a control of the form $\pi^* = \beta\omega$ where ω is a r vector, thus reducing the number of degrees of freedom from n to r, with the goal of investigating the commonly held opinion that an arbitrageur should invest directly in the cointegrating process. Thirdly, we focus exclusively on the time-independent component of the portfolio weights and finally establish its connection to the standard double-threshold strategy.

3.1. The arithmetic Brownian motion asset price model

In the strategy outlined in the previous section, we used a geometric Brownian motion as a building block for modelling the cointegrated prices. Furthermore, since we wanted to derive an analytical solution while using this approach, it was necessary for us to formulate the problem such that the cointegrated time series are the log-prices, rather than the asset prices themselves. Here we assume that the cointegrated processes are linear combinations of the asset prices themselves and we replace the geometric Brownian motion by an arithmetic Brownian motion in the asset prices model. The derivation of the solution of the arithmetic model is almost the same as in the log normal case and yields

$$\tilde{\pi}^*(t,z) = \frac{1}{\gamma}\Bigg[\left(I_n - \beta\delta^t(T-t)\right)(\sigma\sigma^t)^{-1}(\mu+\delta z) - \beta\delta^t(\sigma\sigma^t)^{-1}\delta b\frac{(T-t)^2}{2}\Bigg],$$

where $\tilde{\pi}^*(t,z)$ now denotes the number of shares of these assets held rather than the amount invested. In the remainder of this paper, we name this second type of formulation Normal Asset Model (NAM in short), whereas the first one is referred to as Log Normal Asset Model (LNAM in short).

Exchange	Type	Price	Amount (BTC)	Value ($)	Timestamp
BTC-e	buy	416	3.32	1381.12	Nov 20, 2013, 11:38 AM
BTC-e	withdraw	NA	3.32	NA	Nov 20, 2013, 11:50 AM
Bitstamp	deposit	NA	3.319	NA	Nov 20, 2013, 5:49 PM
Bitstamp	sell	594.97	3.319	1974.71	Nov 20, 2013, 6:09 PM

Figure 1. Example of a buy-move-sell pair trade.

Leg 1:					
Exchange	Type	Price	Amount (BTC)	Value ($)	Timestamp
Bitstamp	Sell	320	0.9515	304.64	Nov 11, 2013, 1:18 AM
BTC-e	Buy	296.197	0.952	281.98	Nov 11, 2013, 1:19 AM
Leg 2:					
Exchange	Type	Price	Amount (BTC)	Value ($)	Timestamp
Bitstamp	Buy	398.98	1.9	758.06	Nov 13, 2013, 10:37 PM
BTC-e	Sell	395.04	2.018	797.2	Nov 13, 2013, 10:24 PM

Figure 2. Example of an instantaneous pair trade.

3.2. *The normal asset model with investment in the cointegrating process*

While the strategies up to this point in the paper have left the investor with the freedom to simply invest in an optimal way, most strategies used in practice have the characteristic that the positions taken are a linear combination of the vectors of cointegrating coefficients. This may sound unfamiliar, but in fact it is a constraint built into many pairs trading strategies and one might ask whether this constraint would have an effect on the optimal trading trajectory. Mathematically, this constraint is expressed as

$$\underline{\pi}^* = \beta\omega,$$

where $\omega \in \mathbb{R}^r$, which represents the vector of weights for each of the r components of the cointegrating process, must be determined. It is easy to see that, under this restriction, the optimal strategy takes the form

$$\underline{\pi}^*(t,z) = \frac{1}{\gamma}\Bigg[\Big(\beta(\beta^t\sigma\sigma^t\beta)^{-1}\beta^t - \beta\delta^t(\sigma\sigma^t)^{-1}(T-t)\Big) \times (\mu+\delta z) - \beta\delta^t(\sigma\sigma^t)^{-1}\delta b\frac{(T-t)^2}{2}\Bigg].$$

It is worth noting that this solution is the orthogonal projection of the unconstrained solution (the NAM strategy) onto the r-dimensional subspace of mean-reverting processes spanned by the columns of β. In the rest of this paper, we use the abbreviation NAMIC for this strategy.

3.3. *Time-independent normal asset model with investment in the cointegrating process*

With the aim of relating our strategy to the double-threshold strategy, we simply truncate the time-dependent term from the previous NAMIC model, leading to the suboptimal strategy

$$\underline{\pi}_0^*(z) = \frac{1}{\gamma}\Big[\beta(\beta^t\sigma\sigma^t\beta)^{-1}\beta^t(\mu+\delta z)\Big].$$

This constitutes a significant simplification that brings the problem closer to familiar territory to most traders. However, as illustrated by the out-of-sample experiments reported in the next section, this simplification will significantly reduce the accumulated profit. In the remainder of this paper, we abbreviate this strategy TINAMIC, the first two letters standing for Time-Independent.

3.4. *Connection to the classic double-threshold strategy*

We show below that the classic double-threshold strategy, which is only applicable in the case when there is a unique cointegration relation ($r = 1$), can be viewed as a simplification or quantization of the TINAMIC strategy. In this subsection, we assume that $r = 1$ and establish this connection.

In the arithmetic model, assuming $b = 0$ to simplify, the long-run mean of the process z is given by

$$\tilde{\nu} = -(\beta^t\delta)^{-1}\beta^t\mu,$$

and the long-run volatility coefficient of z is the constant

$$\sigma_z = \frac{1}{\sqrt{2}}(-\beta^t\delta)^{-\frac{1}{2}}||\sigma^t\beta||.$$

The double-threshold strategy can be then described as follows: one usually sets thresholds above and below the long-run average. The upper bound is of the form $\tilde{\nu} + \lambda\sigma_z$, where the parameter λ is typically set to 2, whereas the lower bound is $\tilde{\nu} - \lambda\sigma_z$. Each time z rises above the threshold $\tilde{\nu} + \lambda\sigma_z$, one takes the position $\underline{\pi}_0 = K\beta$, where K, which dictates the size of the position, is a parameter to be determined, and holds this position for as long as the process z stays above the lower bound. If z falls below the lower bound, then one takes the opposite position $\underline{\pi}_0 = -K\beta$ instead and holds it for as long as z stays below the upper threshold.

Furthermore, we notice that, when z crosses one of the thresholds, the value of the TINAMIC strategy corresponding to that specific value of z equates the double-threshold strategy.

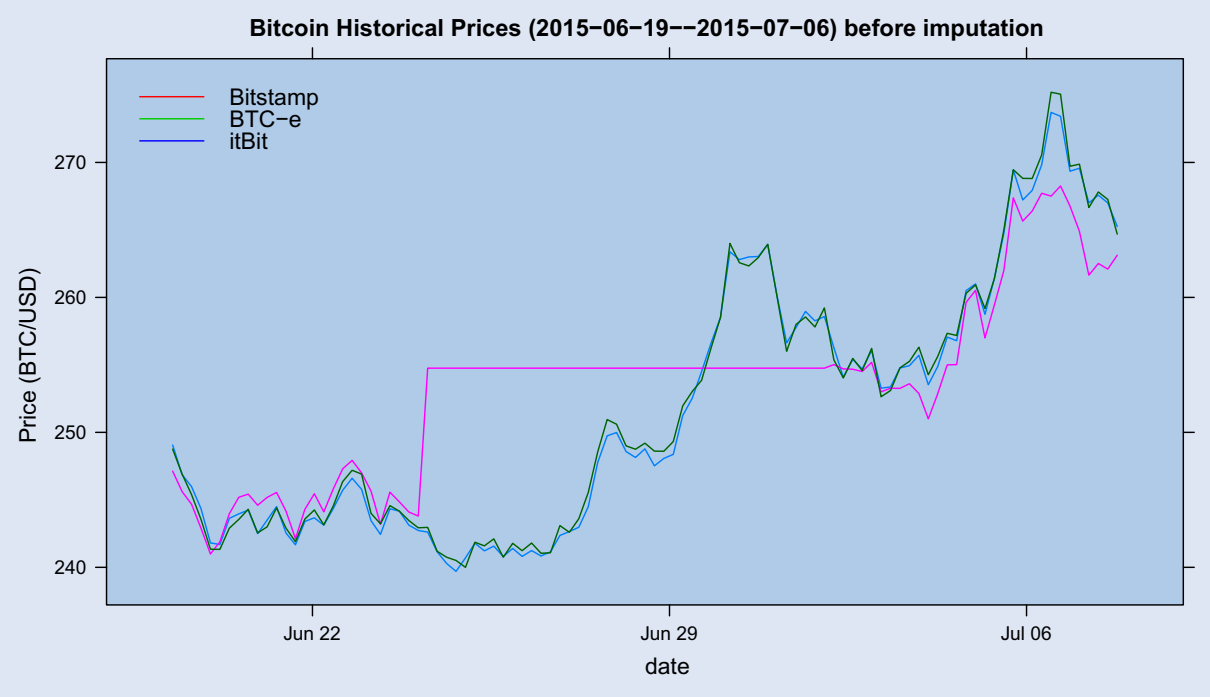

Figure 3. Raw historical price data in us dollars for BTC-e, Bitstamp, and itBit, from 19 June 2015 to 06 July 2015.

In order to see this, we assume for instance that z crosses the upper bound; its value at the time of the threshold-crossing event is then

$$\bar{z} = \tilde{\nu} + \lambda\sigma_z.$$

Substituting $\bar{z}$ into the formula for the TINAMIC control yields

$$\begin{aligned}
\underline{\pi}_0^*(\bar{z}) &= \frac{1}{\gamma}\left(\beta(\beta^t\sigma\sigma^t\beta)^{-1}\beta^t(\mu+\delta\bar{z})\right)\\
&= \frac{1}{\gamma}\left[\beta(\beta^t\sigma\sigma^t\beta)^{-1}\beta^t\left(\mu-\delta(\beta^t\delta)^{-1}\beta^t\mu\right.\right.\\
&\quad \left.\left.+\lambda\delta\frac{1}{\sqrt{2}}(-\beta^t\delta)^{-\frac{1}{2}}||\sigma^t\beta||\right)\right]\\
&= \frac{\lambda}{\sqrt{2}\gamma}(\beta^t\sigma\sigma^t\beta)^{-\frac{1}{2}}(-\beta^t\delta)^{\frac{1}{2}}\beta\\
&= \frac{\lambda}{2\gamma\sigma_z}\beta.
\end{aligned}$$

Finally, we can equate the control computed above to the double-threshold strategy

$$\pi_0^*(\bar{z}) = \underline{\pi}_0 = K\beta.$$

Solving the above equation for K yields

$$K = \frac{1}{2}\frac{\lambda}{\gamma\sigma_z}.$$

We observe that the size of the position is inversely proportional to both the volatility coefficient of the process z and the risk aversion parameter.

Additionally, we include examples of live pair trades from transaction history. In practice, there are two fundamentally different ways of implementing a pairs trade using the double-threshold in bitcoin markets. The more common implementation, which we refer to as *buy-move-sell*, involves buying bitcoins on one exchange, sending them to the other exchange using the blockchain network and then selling them at the destination (see figure 1). This implementation seems attractive to many bitcoin traders because it leverages the essentially free withdrawals and deposits on the blockchain. A direct implementation, or in other words, an *instantaneous pair trade*, is preferable, since there is no transfer period during which the trader is exposed to currency risk. However, it is only feasible when there are dollars and bitcoins present on all exchanges at all times. This in turn means that the trader must have a cheap credit line for bitcoins, which is somewhat difficult to find, but possible using peer-to-peer bitcoin lending markets (see figure 2). While the two legs of the instantaneous pairs trade are not of the same size, this example demonstrates that such spreads occurred with sufficient market depth and frequency to allow for a round-trip pair trade using a double-threshold in two days. By a round trip, we mean two successive pair trades triggered by crossing over the upper threshold and then back under the lower threshold, thereby restoring the capital allocation to its original state.

In the section that follows, we contribute further to the study of the applicability of our model, by backtesting it using one set of historical data in the bitcoin markets.

4. Experiments in the bitcoin markets

We illustrate the results of the previous sections by testing the computed optimal trading strategies in the bitcoin markets. Clearly, this experiment does not constitute a systematic study of the model's performance.

The data used for the test described below are from bitcoincharts.com, and dates from 04 January 2014 to 03 June 2016. In January 2014, a major market disaster caused unprecedented turmoil in the bitcoin markets which, in turn, resulted in extreme spreads and volatility (even by bitcoin standards).

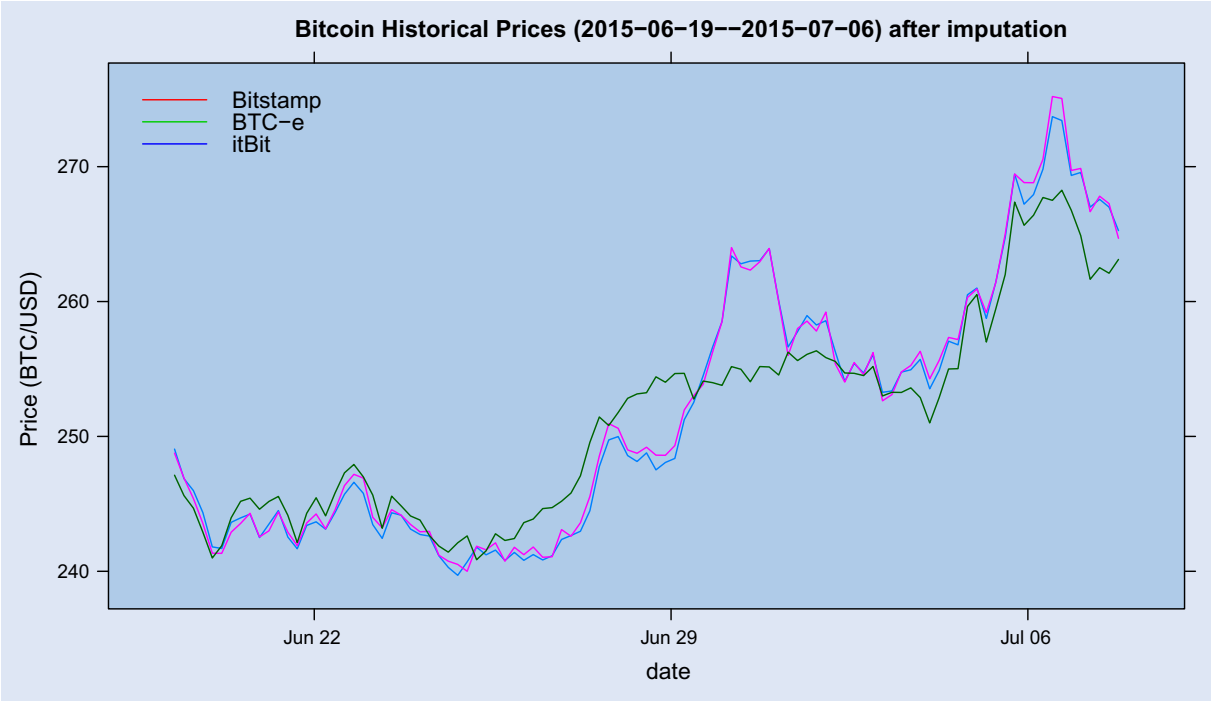

Figure 4. Historical price data in us dollars for BTC-e , Bitstamp and itBit, after imputation with a Brownian bridge.

Table 1. In-sample estimates of model parameters.

Strat	Exchange	Order	β	δ	a	b	μ	γ
LNAM	bitstampUSD	1st order	1.00E+00	−1.59E−01	8.62E−04	1.14E−06	7.92E−05	1.04E+02
LNAM	bitstampUSD	2nd order	0.00E+00	2.17E−02	2.58E−02	−2.41E−06	7.92E−05	1.04E+02
LNAM	btceUSD	1st order	−6.94E−18	2.93E−03	8.62E−04	1.14E−06	7.52E−05	1.04E+02
LNAM	btceUSD	2nd order	1.00E+00	−9.22E−02	2.58E−02	−2.41E−06	7.52E−05	1.04E+02
LNAM	itbitUSD	1st order	−1.00E+00	1.35E−01	8.62E−04	1.14E−06	6.10E−05	1.04E+02
LNAM	itbitUSD	2nd order	−1.00E+00	1.58E−02	2.58E−02	−2.41E−06	6.10E−05	1.04E+02
NAM	bitstampUSD	1st order	1.00E+00	−1.27E−01	−1.77E+00	5.57E−04	−1.68E−02	3.54E+01
NAM	bitstampUSD	2nd order	0.00E+00	2.87E−02	2.28E+00	−9.67E−04	−1.68E−02	3.54E+01
NAM	btceUSD	1st order	0.00E+00	3.86E−02	−1.77E+00	5.57E−04	−1.69E−02	3.54E+01
NAM	btceUSD	2nd order	1.00E+00	−1.09E−01	2.28E+00	−9.67E−04	−1.69E−02	3.54E+01
NAM	itbitUSD	1st order	−9.99E−01	1.46E−01	−1.77E+00	5.57E−04	−1.63E−02	3.54E+01
NAM	itbitUSD	2nd order	−9.91E−01	1.95E−02	2.28E+00	−9.67E−04	−1.63E−02	3.54E+01
NAMIC	bitstampUSD	1st order	1.00E+00	−1.27E−01	−1.77E+00	5.57E−04	−1.68E−02	3.54E+01
NAMIC	bitstampUSD	2nd order	0.00E+00	2.87E−02	2.28E+00	−9.67E−04	−1.68E−02	3.54E+01
NAMIC	btceUSD	1st order	0.00E+00	3.86E−02	−1.77E+00	5.57E−04	−1.69E−02	3.54E+01
NAMIC	btceUSD	2nd order	1.00E+00	−1.09E−01	2.28E+00	−9.67E−04	−1.69E−02	3.54E+01
NAMIC	itbitUSD	1st order	−9.99E−01	1.46E−01	−1.77E+00	5.57E−04	−1.63E−02	3.54E+01
NAMIC	itbitUSD	2nd order	−9.91E−01	1.95E−02	2.28E+00	−9.67E−04	−1.63E−02	3.54E+01
TINAMIC	bitstampUSD	1st order	1.00E+00	−1.27E−01	−1.77E+00	5.57E−04	−1.68E−02	6.76E−02
TINAMIC	bitstampUSD	2nd order	0.00E+00	2.87E−02	2.28E+00	−9.67E−04	−1.68E−02	6.76E−02
TINAMIC	btceUSD	1st order	0.00E+00	3.86E−02	−1.77E+00	5.57E−04	−1.69E−02	6.76E−02
TINAMIC	btceUSD	2nd order	1.00E+00	−1.09E−01	2.28E+00	−9.67E−04	−1.69E−02	6.76E−02
TINAMIC	itbitUSD	1st order	−9.99E−01	1.46E−01	−1.77E+00	5.57E−04	−1.63E−02	6.76E−02
TINAMIC	itbitUSD	2nd order	−9.91E−01	1.95E−02	2.28E+00	−9.67E−04	−1.63E−02	6.76E−02

The markets during this period were also highly illiquid, and therefore since extremely high transaction costs would likely distort results, this period and any history before then were excluded from the analysis. The trade prices in US dollars on the three exchanges BTC-e, Bitstamp and itBit are downloaded from this website.

Next, we have to select a sensible sampling frequency accounting for the presence of microstructure noise in the data. Note that the microstructure noise was not incorporated in the stochastic problem proposed in section 2 and that we simply assumed that the variables x representing the log-prices were observable. First of all, our raw data-set, which contains all the trade prices, exhibits some significant serial autocorrelations. These become insignificant if the time interval between two adjacent samples is greater or equal to 1000 s. Besides, we apply formula (19) in Ait-Sahalia *et al*. (2005) providing the optimal sampling frequency of an asset, given an estimate of the standard variation of the microstructure noise, of the asset

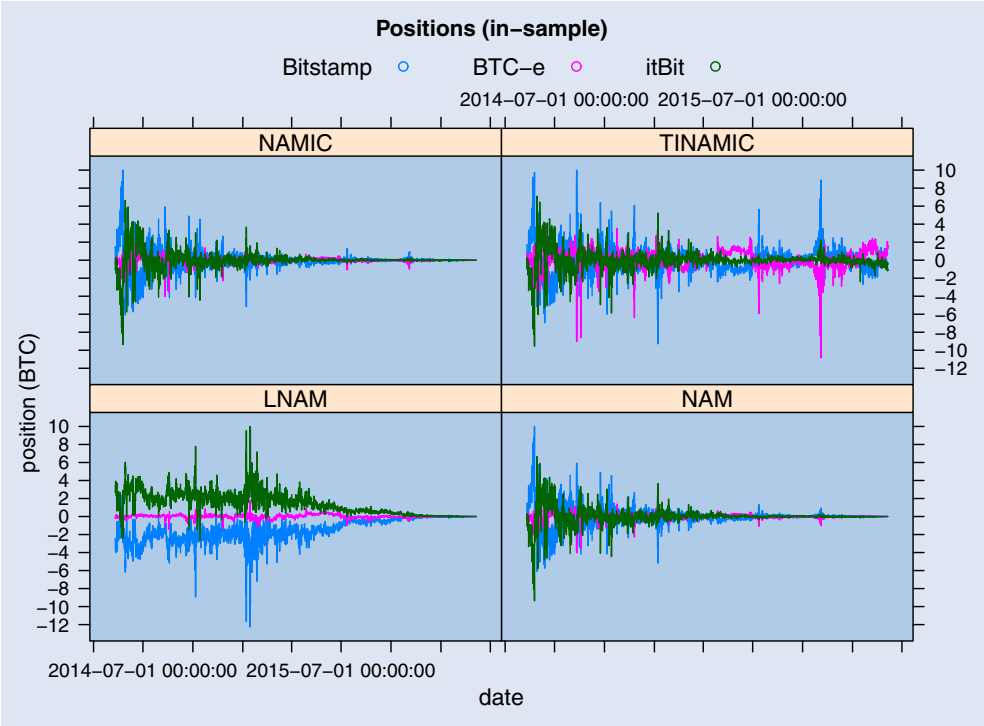

Figure 5. The positions (in BTC) in the three exchanges over time for the four trading strategies in an in-sample test. Note that the positions for all strategies are normalized so that their maximum positions over the testing interval are equal.

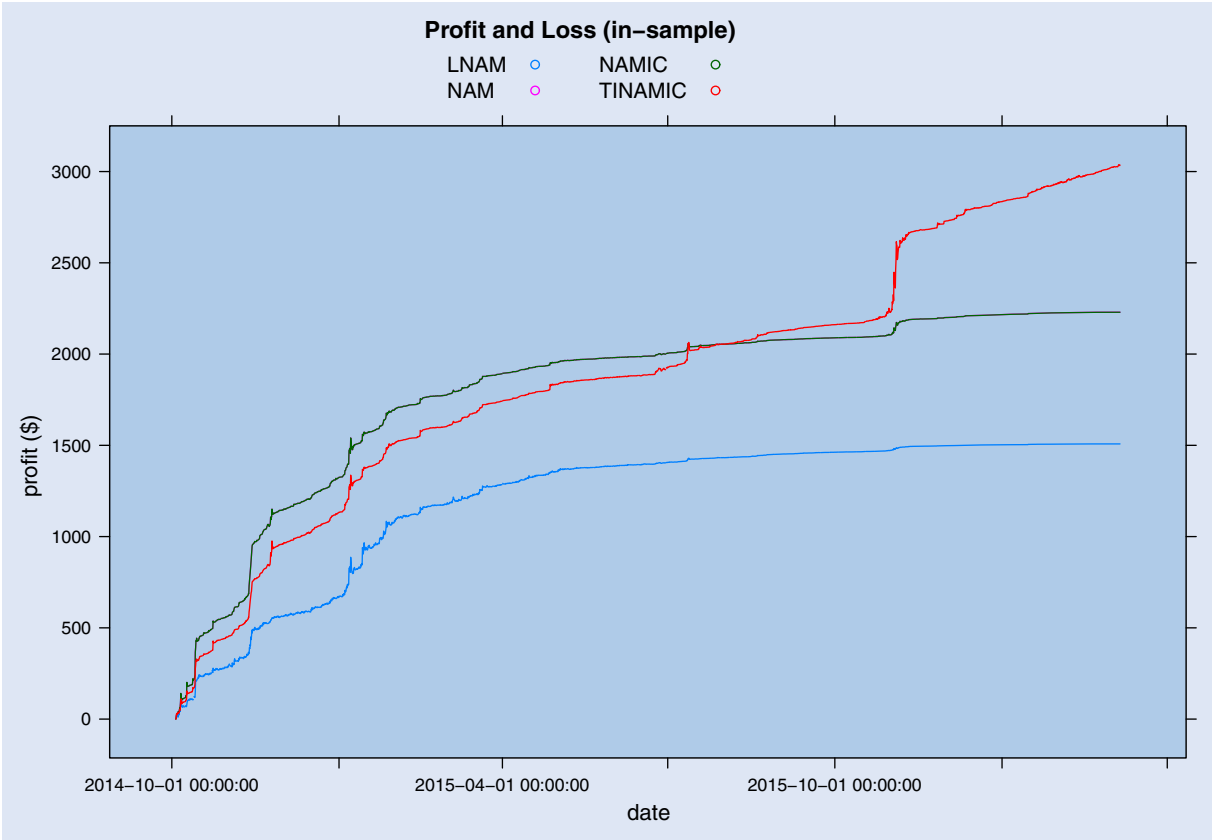

Figure 6. Profit and Loss (PnL) generated by in-sample backtesting of the four strategies. Note that the NAM trajectory is barely visible due to its overlap with the NAMIC trajectory.

volatility and of the length of the time-interval used for estimating the volatility. To this end, we compute an estimate of the standard deviation of the microstructure noise a in the following manner: to simplify, we assume that the bid-ask spread s is the sole source of noise and set a to $s/2$. On the three exchanges we consider, the bid-ask spread rarely reaches up to 2% of the asset price. Considering this as a three sigma event yields the estimate $a = 0.3\%$. Given that the annualized volatility is

Table 2. In-sample estimate of the covariance matrix of the log-returns for 16, 000 s intervals (main model).

	Bitstamp	BTC-e	itBit
Bitstamp	0.0002363	0.0002067	0.0001868
BTC-e	0.0002067	0.0002275	0.0001717
itBit	0.0001868	0.0001717	0.0001979

about 0.7, and using a 40-day rolling window for the estimation, (19) yields an interval of about 16 000 s. Consequently, the data we use in the tests we present below are sampled at 16 000 s intervals. It is worth mentioning that we also conducted additional out-of-sample experiments at various sampling frequencies, namely for 1000, 5000, 8000, 16 0000, 40 000 s intervals, and although, the results differed somewhat, they all yielded a positive profit.

When dealing with any markets that are inefficient enough to apply this kind of analysis, it is practical to expect that there may be limitations in the quality of the data. The data from bitcoincharts had a major issue: some of the exchanges would have occasional extended periods with no price change, whether because of technical problems, extended Distributed Denial of Service attacks or simply because of low volume. This has a drastic effect on the spreads of the exchanges, and can cause spikes in the estimates of the matrix β, which are then amplified by spikes in the process z, thus creating large discontinuities in the optimal trading trajectory. In order to solve this problem, the time series for each exchange were processed through a filter that identified extended periods without a price change and then imputed these areas with a series of Brownian bridges. In the data-set used, there were only two such extended periods without any price change, with the longest one on Bitstamp lasting one week (see figures 3 and 4 below).

The presence of cointegration was tested using the Johansen test implemented in the R library tsDyn. The cointegration matrices δ and β^t, as well as the parameters a, b in the cointegrating relations, were also estimated with the package tsDyn. The columns of β are ordered in decreasing magnitude of their corresponding eigenvalue. So the first element of z, associated with the first column of β, also corresponds to the largest eigenvalue, while the second element of z corresponds to the second largest eigenvalue. In the discussion and the figures that follow, we refer to the elements of z as first- and second-order components of z. Finally, the other parameters in the model were estimated using the maximum likelihood method.

We implemented the four strategies derived from our model, namely LNAM, NAM, NAMIC and TINAMIC, and conducted both in-sample and out-of-sample tests for all of these. This necessitates estimating all the parameters in the model numerous times, once for the in-sample test on the whole time frame, and in all the successive time windows for the out-of-sample test. For the sake of brevity, we only show here the estimated in-sample parameters in tables 1 and 2. We observe that the value of b turns out to be very small for this data-set when the whole time interval is used. For the out-of-sample tests, we update the parameters regularly using a 40-day look-back window. As explained above, the choice of the sampling frequency is consistent with a length of 40 days for the sliding window.

Since we are using the exponential utility function for convenience, the initial capital is irrelevant as it factors out of the value function. In other words, the initial wealth is arbitrary and the rate of return cannot be computed. For instance, the initial capital can be simply taken equal to 0. For each application of the model, we compute here the resulting cumulative profit and loss in US Dollars.

Besides, as we can see in (14)–(17), the size of each position is inversely proportional to the risk aversion parameter γ. In order to standardize the graphs among the various strategies, the values of γ were chosen such that the maximum position for all of the strategies over the displayed time period is the same. While this choice is important for visualization purposes, we also take into account the average daily trading volume of the strategy when choosing γ. Thus, the maximum position was set to 10 BTC so that the average daily trading volume by all of our models stays below 0.1% of the average daily volume of the least liquid market. In summary, the strategies are first run in order to determine their maximum position over the time period, and then the values of γ, position sizes, and PnL curves are adjusted so that these constraints on visualization and market impact costs are reflected. For instance, for the in-sample test for the LNAM strategy to achieve a maximum position of 10 BTC yields a value of γ equal to 103.59 as seen in table 1. Table 3 shows that for the LNAM strategy corresponding to this choice of γ, the average daily trading volume on Bitstamp, BTC-e and itBit are, respectively, equal to 2.73 BTC, 0.465 BTC and 2.42 BTC. Note that the maximum daily trading volume over the testing period got as high as 31.9 BTC on Bitstamp, which is roughly 1% of average daily volume on that exchange. In practice, the calibration of γ would likely be done iteratively using a running optimization, but the appropriate implementation of this calibration is a much larger topic and beyond the scope of this paper.

Next, in the out-of-sample test, we apply the model in a consecutive series of intervals, rebalancing our positions with the same frequency as the chosen sampling frequency, that is, every 16 000 s. In each application, the parameters are set to the estimates obtained for the data in the previous window, the initial time is set to the beginning of the current interval and the final time horizon is taken to be the last date in the data-set. Note that the first 40-day window is only used for estimating the first set of parameters and that we do not start trading until the beginning of the second window. Since all of the out-of-sample tests are disjoint from each other in time, we depict the results in a single chart.

Finally, we subtract an estimated transaction fee and bid-ask slippage from the profit and loss. We assume that the transaction fee is proportional to the volume traded. We show in table 4, the numerical values of the constant proportional transaction fees and the bid-ask spread expressed as a percentage of the price for the three exchanges.

Figures 5 and 6 are graphs describing the in-sample performance of the four strategies, so that all market parameters are estimated over the entire period and held constant. Figures 7 and 8 shows the in-sample process z and matrix of cointegration coefficients β for the NAM strategy. The out-of-sample positions and cumulative profits and losses for each of the four strategies are given in figures 9 and 10. The out-of-sample process z and the matrix of cointegration parameters β for the NAM strategy are given, respectively, in figures 11 and 12.

We observe that the in-sample and out-of-sample tests yield similar results. Indeed, in both in-sample and out-of-sample

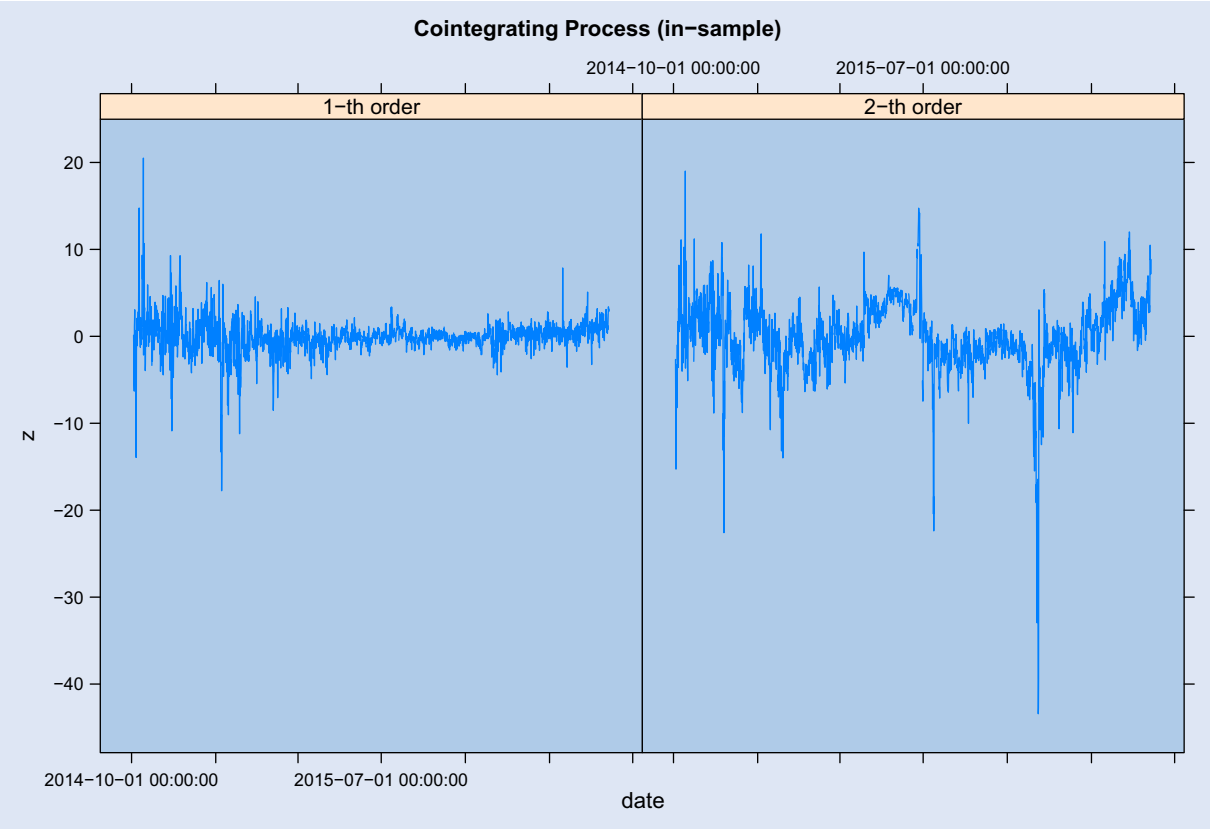

Figure 7. The first- and second-order components of z over time, estimated in-sample, for the arithmetic model.

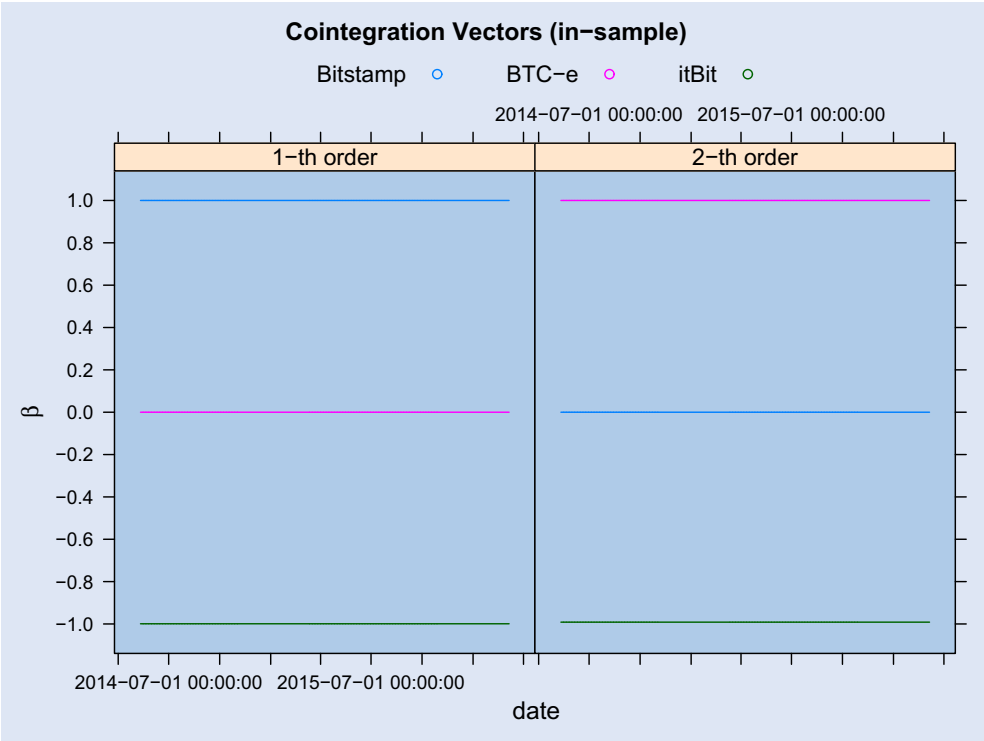

Figure 8. The first- and second-order cointegration vectors over time, estimated in-sample (columns of β), for the arithmetic model. Note that since we estimate the parameters only once for the entire sample, these vectors are constant over time.

tests, the trading trajectories follow either a pattern that is either roughly decreasing (LNAM, NAM and NAMIC), or constant (TINAMIC) (see figures 5 and 9). This is consistent with the fact that the TINAMIC strategy is the only one whose positions do not increase with time to maturity. One consequence of this is that while the TINAMIC strategy performs poorly to

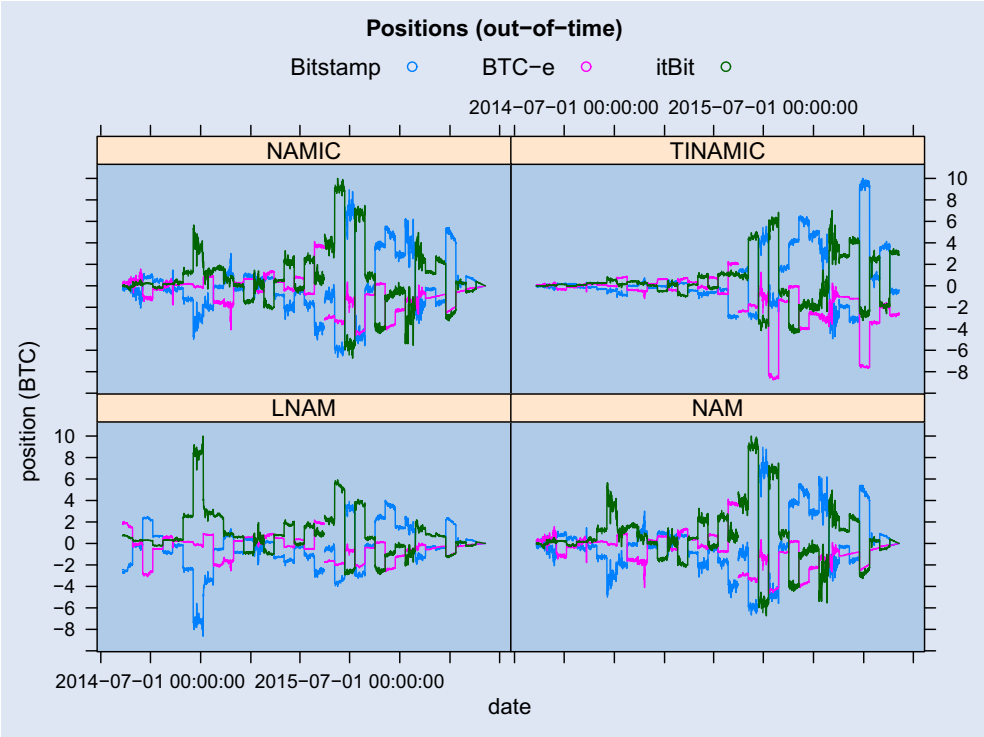

Figure 9. The positions (in BTC) in the three exchanges over time for the four strategies in an out-of-time test. Note that since the positions are generated only using a 40-day rolling window of historical data, we are not able to include the first 40 days in the tests as this window is used for the initial estimation of parameters.

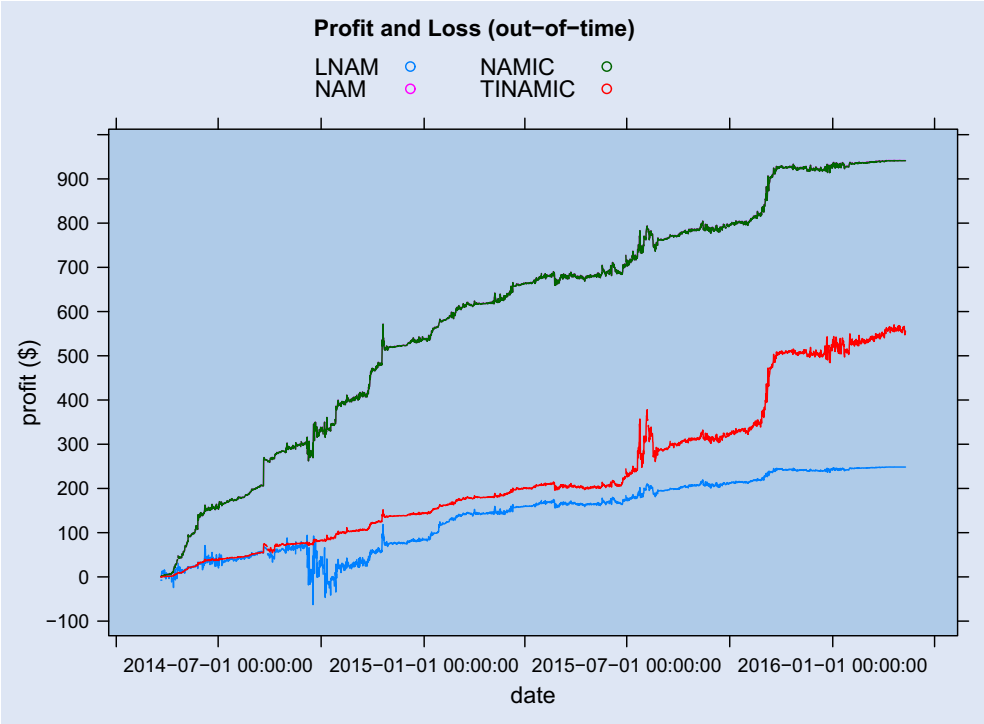

Figure 10. Profit and Loss (PnL) generated by out-of-time backtesting of the four strategies.

moderately towards the beginning of the trading period, it tends to outperform the other strategies towards the end, since it maintains positions of high magnitude even as the time approaches maturity, ignoring the fact that the prices may not revert back before the end of the trading period. Conversely, the other strategies have PnL curves that flatten as they approach maturity, because their positions shrink in magnitude as they approach the time horizon (see figures 6 and 10). In addition, it is clear that the NAM and NAMIC strategies are nearly (though not exactly) identical. This supports the intuition that the investor can make the simplifying assumption that she should consider the cointegrating processes to be the assets

Table 3. Trading volume for each strategy on the three exchanges in BTC.

strategy	exchange	maxDailyTrade	avgDailyTrade	biggestTrade	avgTrade
LNAM	Bitstamp	3.19E+01	2.73E+00	8.25E+00	1.90E−01
LNAM	BTC-e	5.11E+00	4.65E−01	1.80E+00	3.30E−02
LNAM	itBit	2.87E+01	2.42E+00	6.01E+00	1.68E−01
NAM	Bitstamp	3.83E+01	2.76E+00	1.01E+01	1.81E−01
NAM	BTC-e	1.11E+01	8.55E−01	4.52E+00	5.73E−02
NAM	itBit	3.28E+01	2.29E+00	9.54E+00	1.48E−01
NAMIC	Bitstamp	3.83E+01	2.76E+00	1.01E+01	1.81E−01
NAMIC	BTC-e	1.11E+01	8.55E−01	4.52E+00	5.73E−02
NAMIC	itBit	3.28E+01	2.29E+00	9.54E+00	1.48E−01
TINAMIC	Bitstamp	5.46E+01	4.59E+00	1.03E+01	3.18E−01
TINAMIC	BTC-e	3.18E+01	2.83E+00	1.09E+01	1.99E−01
TINAMIC	itBit	4.18E+01	3.20E+00	9.96E+00	2.15E−01

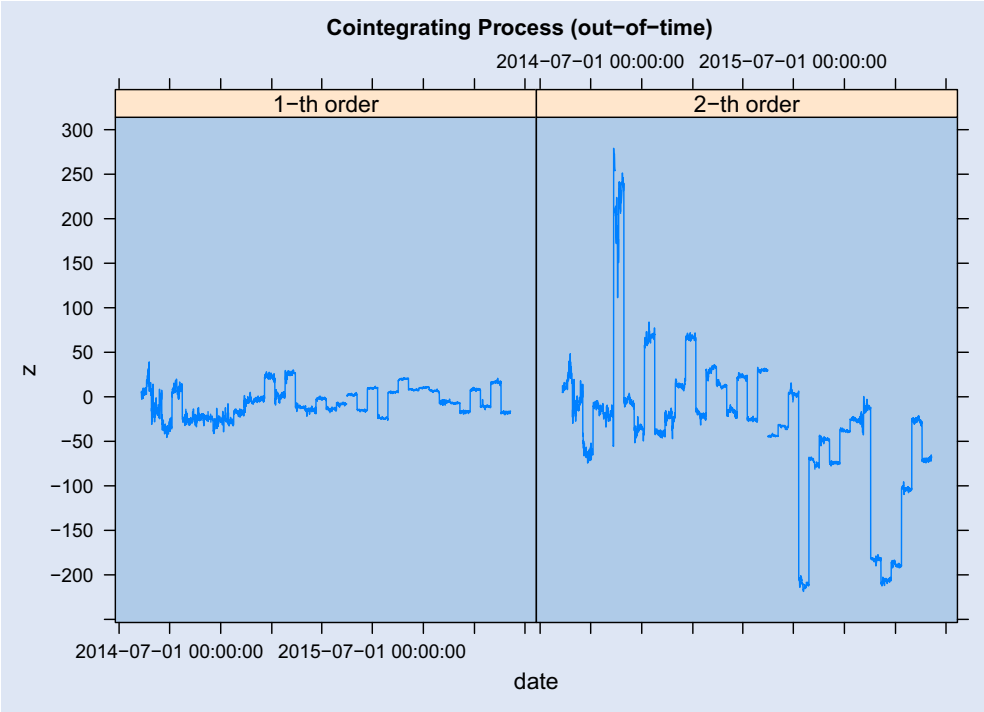

Figure 11. The first- and second-order components of z over time, using a rolling out-of-time window for estimation, for the arithmetic model.

Table 4. Transaction fee and bid-ask spread.

	Bitstamp	BTC-e	itBit
Transaction fees (per number of Bitcoins traded)	0.0025	0.002	0.002
Bid-ask spread (as a fraction of the price in USD)	0.005	0.005	0.01

unto themselves, so that all positions are linear combinations of these cointegration vectors. In the case of two assets with a cointegrating vector close to $(-1, 1)$, this intuition means that the investor can and should *trade the spread*.

5. Conclusion

We proposed a pairs trading model for a portfolio of cointegrated assets, derived the optimal trading strategies in closed-form and proved a verification result. We tested the effects of assuming normal asset dynamics, constraining the positions to be linear combinations of the cointegrating processes and truncating the terms that depend explicitly on the time-to-maturity. We then draw a mathematical connection to the classical double-threshold strategy and spend some time discussing the implementation of such a strategy, using live trading examples. Finally, in order to better understand the model, we performed an out-of-sample test in the bitcoin markets.

While pairs trading in bitcoin markets have certainly been possible historically, this paper is in no way suggesting that the reader go out and implement one of these strategies today using a large amount of capital. Despite the care taken to avoid certain pitfalls in backtesting such as bid/ask slippage, transaction

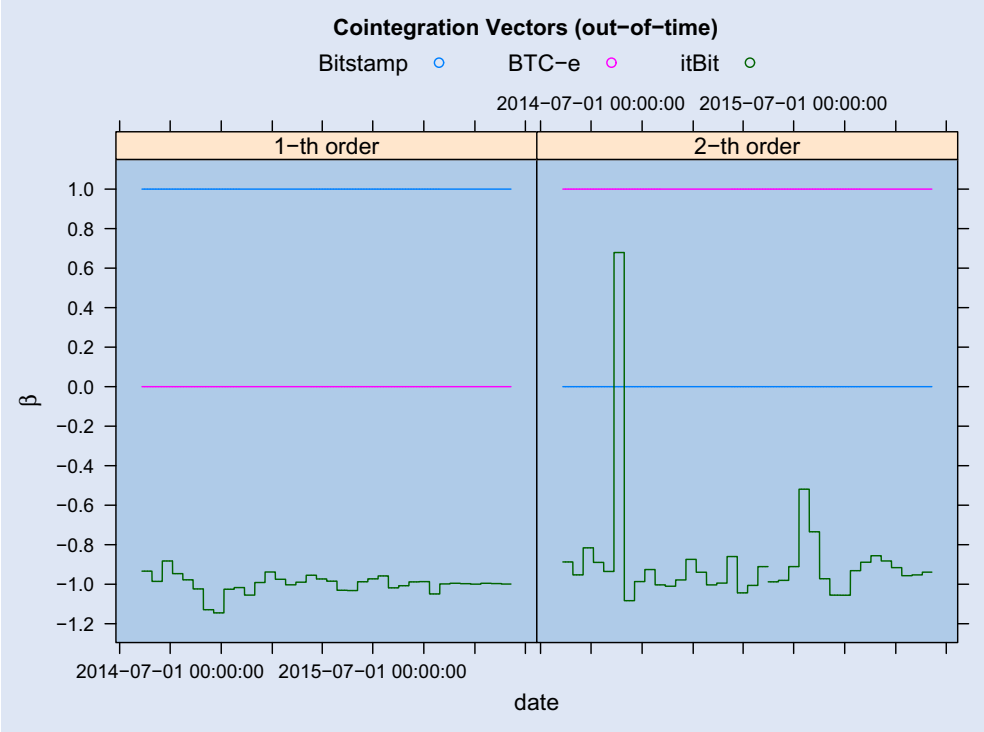

Figure 12. The first- and second-order cointegration vectors over time, using a rolling out-of-time window for estimation, for the arithmetic model.

costs and market impact costs, there are still some significant assumptions being made that could distort the tests results. There are major operational risks involved in implementing an automated trading strategy, including but not limited to latency risk. Secondly, there are times when limited liquidity and market depth could cause significant market impact costs to the trader, even using the limits on maximum position that we have deemed appropriate, and especially using much more capital. These effects would particularly distort the profit and loss at times when the cointegrating process has a very large magnitude, as the liquidity at these times is likely to be low, and the price is likely to be moving very fast, which could easily turn a profitable trade during testing into a real loss.

Notably, the LNAM model has a quadratic dependence on the time-to-maturity because of the convexity adjustment in the lognormal asset dynamics. However, the three strategies that are exact solutions to stochastic control problems—LNAM, NAM and NAMIC—all have at least a linear dependence on the time-to-maturity, which is a more fundamental feature of the problem's formulation. As discussed in section 3, we cannot remedy this dependence by simply taking the limiting case of an infinite horizon, since it causes the value function to become unbounded. In the future, we are planning to investigate the applicability of the forward performance criterion developed in Musiela and Zariphopoulou (2008) to our pairs trading model to eliminate the need to fix the time horizon.

Acknowledgements

We thank Robert Almgren and Peter Carr for their insightful comments and suggestions. We are also very grateful to the editor and the anonymous reviewer for their appropriate and constructive suggestions to improve the paper.

Disclosure statement

No potential conflict of interest was reported by the authors.

References

Ait-Sahalia, Y., Mykland, P.A. and Zhang, L., How often to sample a continuous-time process in the presence of market microstructure noise. *Rev. Finance Stud.*, 2005, **18**(2), 351–416.

Avellaneda, M. and Lee, J.-H., Statistical arbitrage in the U.S. equities market. *Quant. Finance*, 2010, **10**(2), 761–782.

Benth, F.E. and Karlsen, K.H., A note on Merton's portfolio selection problem for the schwartz mean-reversion model. *Stochastic Anal. Appl.*, 2005, **23**, 687–704.

Cartea, A. and Jaimungal, S., Algorithmic trading of co-integrated assets, *Int. J. Theor. Appl. Financ*, forthcoming.

Chiu, M.C. and Wong, H.Y., Mean−variance portfolio selection of cointegrated assets. *J. Econ. Dyn. Control*, 2011, **35**, 1369–1385.

Duan, J.-C. and Pliska, S., Option valuation with co-integrated asset prices. *J. Econ. Dyn. Control*, 2004, **28**, 727–754.

Fleming, W.H. and Soner, H.M., *Controlled Markov Processes and Viscosity Solutions*, 1993 (Springer-Verlag: New York, NY).

Lei, Y. and Xu, J., Costly arbitrage through pairs trading. *J. Econ. Dyn. Control*, 2015, **56**, 1–19.

Leung, T. and Li, X., Optimal mean reversion trading with transaction costs and stop-loss exit. *Int. J. Theor. Appl. Finance*, 2015, **18**(2), 1550020.

Merton, R., Optimum consumption and portfolio rules in a continuous time model. *J. Econ. Theor.*, 1971, **3**, 373–413.

Mudchanatongsuk, S., Primbs, J.A. and Wong, W., Optimal pairs trading: A stochastic control approach. In *Proceedings of the American Control Conference*, Westin Seattle Hotel, Seattle, WA, pp. 1035–1039, 2008.

Musiela, M. and Zariphopoulou, P., Optimal asset allocation under forward exponential performance criteria. In *Markov Processes and Related Topics: A Festschrift for Thomas G. Kurtz*, edited by S. N. Ethier, J. Feng, and R. H. Stockbridge, pp. 285–300, 2008 (Institute of Mathematical Statistics: Beachwood, OH, USA).

Ngo, M.-M. and Pham, H., Optimal switching for pairs trading rule: A viscosity solutions approach, *J. Math. Anal. Appl.*, 2016, **441**(1), 403–425.

Pham, H., *Continuous-time Stochastic Control and Optimization with Financial Applications: Stochastic Modelling and Applied Probability*, 2009 (Springer: Berlin, Heidelberg).

Tourin, A. and Yan, R., Dynamic pairs trading using the stochastic control approach. *J. Econ. Dyn. Control*, 2013, **37**(10), 1972–1981.