

A Two-Factor Model of Stock Returns and its Application in Pairs Trading

1. Introduction

Momentum and mean reversion are two fundamental forces that shape the dynamics of asset prices in financial markets. Traditional financial models often struggle to accurately capture these complex behaviors, typically focusing on either aspect in isolation. This paper introduces an innovative two-factor model that bridges this gap, offering a comprehensive framework that effectively incorporates both momentum and mean reversion effects within the stochastic processes governing financial assets.

The proposed model represents a significant advancement in financial econometrics, providing a nuanced lens through which the multifaceted nature of market behaviors can be understood and predicted. The motivation for this model stems from the empirical observation that financial markets exhibit distinct phases of persistent trends (momentum) and reversion to historical means or intrinsic values (mean reversion). Capturing this duality within a single framework has been a long-standing challenge in financial modeling, which this research aims to address.

The two-factor model innovatively combines a drift component, embodying the momentum effect, with a mean-reverting process, encapsulating the tendency of prices to revert to a long-term equilibrium. The first factor, $d\mu(t)$, represents the long-term trend or drift of an asset's price, integrating the momentum effect through a constant drift parameter $\bar{\theta}$. This component reflects the underlying momentum driven by broader market forces or fundamental changes. The second factor, $d\theta_t$, models the short-term deviations from the drift through its mean-reversion speed κ , capturing the asset's propensity to revert to its mean following short-term fluctuations. These components are modulated by their respective volatilities ($\sigma_\mu, \sigma_\theta$) and driven by correlated Wiener processes, enabling the model to reflect the real-world interplay between momentum and mean reversion.

To demonstrate the model's applicability and effectiveness, this research applies the two-factor model to daily returns data of Coca-Cola (KO) and PepsiCo (PEP) over a twenty-year period, illustrating its potential for informing pairs trading strategies. The parameter estimation employs a sophisticated maximum likelihood estimation (MLE) technique, tailored to accommodate the intricacies of fitting a two-factor model to empirical data, ensuring both accuracy and adaptability in evolving market conditions.

This paper contributes a novel perspective to the quantitative finance literature by presenting a model that integrates key market behaviors—momentum and mean reversion—within a unified analytical framework. By doing so, it offers profound insights into asset price dynamics, enhancing the predictive capabilities of financial models and informing more effective trading strategies. The findings have significant implications for both academic research and practical applications in financial markets, positioning this model as a pivotal tool for investors, analysts, and policymakers.

The remainder of this paper is structured as follows: Section 2 provides a comprehensive literature review, highlighting the existing research on momentum and mean reversion in financial markets. Section 3 introduces the proposed two-factor model, detailing its mathematical formulation and theoretical underpinnings. Section 4 describes the data and methodology employed in the empirical analysis, including the MLE approach for parameter estimation. Section 5 presents the results and discusses the implications of the model for pairs trading strategies. Finally, Section 6 concludes the paper, summarizing the key findings and outlining potential avenues for future research.

2. Literature Review

The development of two - factor models in finance has been largely motivated by the need to capture the complex dynamics of asset prices, particularly in the context of commodity and energy markets . This literature review focuses on key contributions that have shaped the understanding and application of two - factor models in these domains.

Longstaff and Schwartz (1992) develop a two-factor general equilibrium model for interest rate volatility and the term structure. Their model incorporates both a short-term and a long-term factor, providing a framework for analyzing the behavior of interest rates across different maturities.

One of the seminal works in this area is the paper by Schwartz (1997), which introduces a two - factor model for commodity prices . The model incorporates both short - term deviations and long - term equilibrium dynamics, providing a framework for analyzing the stochastic behavior of commodity prices . Schwartz demonstrates the model' s implications for valuation and hedging, highlighting its practical relevance for market participants.

Building upon this foundation, Schwartz and Smith (2000) propose a two - factor model that explicitly separates short - term variations from long - term dynamics in commodity prices . Their model includes a mean - reverting factor for short - term deviations and a Gaussian factor for long - term equilibrium prices.

This separation allows for a more nuanced analysis of commodity price behavior and has been widely influential in subsequent research.

Ribeiro and Hodges (2004) propose a two - factor model specifically designed for commodity prices and futures valuation. Their model incorporates a stochastic convenience yield, which represents the benefit of holding the physical commodity, and a stochastic interest rate. The authors show that their model can effectively capture the dynamics of commodity futures prices and provide insights into the valuation of futures contracts.

Casassus and Collin - Dufresne (2005) extend the two - factor framework by introducing a stochas-

tic convenience yield implied from commodity futures and interest rates . Their model allows for a time - varying convenience yield, which is a key determinant of commodity futures prices. The authors demonstrate the model' s ability to capture the joint dynamics of commodity spot prices, futures prices, and interest rates, providing a comprehensive framework for analyzing commodity markets.

In equity markets, Adrian and Rosenberg (2008) develop a two-factor model for stock returns and volatility, where the factors represent the short-run and long-run components of market risk. They show that this decomposition of market risk helps explain the time-series and cross-sectional variation in stock returns and has implications for asset pricing and risk management.

These papers represent key milestones in the development and application of two - factor models in commodity, equity and energy markets. They highlight the importance of capturing both short - term variations and long - term dynamics, as well as incorporating market - specific factors such as convenience yields and interest rates . While primarily focused on commodities and energy, the insights gained from these models have laid the groundwork for the extension of two - factor approaches to other financial assets, including equities. The proposed two - factor equity model aims to build upon this rich literature by adapting and extending the two - factor framework to capture the unique characteristics of equity markets. By incorporating a drift component for momentum and a mean - reverting component for long - term equilibrium dynamics, the proposed model seeks to provide a comprehensive framework for analyzing equity price behavior, drawing inspiration from the seminal works in commodity and energy markets .

3. Formulation of the Two - Factor Model

The model is defined by a set of SDEs that describe the evolution of an asset' s price, integrating both the long-term drift and the short-term mean-reverting behavior, each subjected to random shocks :

Long - term Drift Component (θ) :

The long - term drift or trend in the asset' s price is modeled as :

$$d\mu_t = \theta_t dt + \sigma_\mu dW_t \quad (1)$$

- $d\mu_t$ represents the change in the long - term component of the asset' s price.
- θ_t is the time-varying drift parameter, indicating the direction and magnitude of the trend.
- σ_μ denotes the volatility associated with the long - term drift.
- dW_t is a Wiener process representing the random fluctuations impacting the long - term trend.

Short - term Mean - Reverting Component :

The short - term deviations from the long - term trend are captured by a mean - reverting process :

$$d\theta_t = -\kappa [\theta_t - \bar{\theta}] dt + \sigma_\theta dZ_t \quad (2)$$

- $d\theta_t$ represents the change in the short - term deviation from the long - term mean.
- κ is the speed of mean reversion, indicating how quickly deviations from the long - term trend are corrected.
- σ_θ denotes the volatility associated with the short - term deviations.
- dZ_t is another Wiener process, independent from dW_t , representing the random shocks to the short-term deviations.

Correlation Between Stochastic Processes

The model allows for a correlation between the two stochastic processes dW_t and dZ_t denoted by ρ . This term captures the interdependence between the long - term and short - term dynamics of asset prices.

Implications of the Model

The inclusion of both stochastic differential equations in the model allows for a dynamic representation of asset prices, accommodating mean-reverting fluctuations around a long - term trend while accounting for the market's inherent uncertainty and noise. The model's structure facilitates the analysis of how fundamental trends and market sentiment drive asset prices over time, alongside the impact of short - term market volatilities and corrections .

4. Positioning within the Landscape of Financial Models

To better understand the significance of the proposed two-factor model, it is useful to compare it with well-established financial models such as the Geometric Brownian Motion (GBM) model and the Ornstein-Uhlenbeck (O-U) model. This comparison highlights the unique features of the two-factor model and its potential to bridge the gap between trend-following and mean-reverting behaviors in asset price dynamics.

The classical GBM model, defined as $dS_t = \mu S_t dt + \sigma S_t dB_t$, shares some fundamental characteristics with the two-factor model. Both models incorporate stochastic processes to capture the inherent randomness in asset price movements. The drift term (μ) in the GBM model is conceptually similar to the long-term component (θ) in the two-factor model, representing the expected return or trend of the asset over time. However, the GBM model assumes that prices follow a log-normal distribution without reverting to a long-term mean, making it more suitable for modeling assets that exhibit indefinite growth, such as stock prices in the long run.

In contrast, the O-U model, defined as $dX_t = \kappa(\theta - X_t)dt + \sigma dB_t$, is a mean-reverting stochastic process that focuses on the tendency of a process to revert to a mean level (θ) over time, with a speed of mean reversion (κ). The mean-reverting component of the two-factor model shares this fundamental characteristic with the O-U model. However, the O-U model is a single-factor model that solely addresses mean reversion, while the two-factor model combines both mean reversion and a separate drift component, offering greater flexibility in modeling complex asset price behaviors.

The two-factor model's ability to capture both short-term deviations and long-term trends sets it apart from models that focus on either trend-following or mean-reverting behavior in isolation. This dual approach allows for a more comprehensive analysis of asset price dynamics, making it applicable to a wide range of financial instruments, including stocks, commodities, and currencies.

Other relevant models, such as the Vasicek model for interest rates and the Heston model for stochastic volatility, introduce additional layers of complexity in financial modeling. While the Vasicek model is similar to the O-U model in its focus on mean reversion, it is specifically designed for interest rate dynamics. The Heston model, on the other hand, incorporates stochastic volatility into option pricing, differing from the fixed volatility assumption in GBM and the simple mean-reversion in the O-U model. The Schwartz two-factor model, introduced by Eduardo Schwartz in 1997, is a well-established model in the field of commodity pricing and has been widely applied to various financial markets. However, in the Schwartz model, the two factors represent the short-term deviation from the long-term equilibrium price and the long-term equilibrium price itself. The equity two-factor model, on the other hand, interprets the two factors as a drift component capturing momentum and a mean-reverting component capturing the tendency to revert to a long-term mean.

The two-factor model's unique contribution lies in its ability to integrate both momentum and mean reversion within a single framework. By combining the long-term drift component (θ) with the mean-reverting component (κ), the model provides a more nuanced perspective on asset price dynamics, capturing the complex interplay between trend-following and mean-reverting behaviors.

5. Model Estimation

The basic procedure for estimating the model is, firstly, to convert the continuous-time SDEs to a discrete approximation using a numerical scheme such as Euler-Maruyama or Milstein. This step is crucial for practical implementation since it allows for the use of observed data in discrete time intervals for parameter estimation and model testing. The next step involves deriving the log-likelihood function based on the discretized model under the assumption of Normality. Finally, with the log-likelihood function in hand, use numerical optimization techniques to find the set of parameters (θ , σ_θ , κ , σ_χ and ρ) that maximize this function. This step is computationally intensive and requires careful selection of optimization algorithms to ensure convergence and accuracy.

Discretization and Log - Likelihood

The Euler-Maruyama method offers a straightforward yet powerful tool for approximating and analyzing complex dynamics described by SDEs, enabling the practical application of sophisticated models like the proposed two - factor framework in various financial contexts.

We discretize the SDEs using the Euler-Maruyama method for a time step from t to $t + \Delta t$, as follows:

$$\mu_{t+\Delta t} = \mu_t + \theta_t \Delta t + \sigma_\mu \Delta W_t \quad (3)$$

- $\mu_{t+\Delta t}$ and μ_t represent the values of the μ components at times $t+\Delta t$ and t respectively.
- ΔW_t is the increment of the Wiener process W_t over the interval Δt , typically simulated as $\sqrt{\Delta t} \times N(0,1)$, where $N(0,1)$ is a standard normal random variable.

Similarly, for $d\theta_t$, the approximation over the same time step is:

$$\theta_{t+\Delta t} = \theta_t - \kappa (\theta_t - \bar{\theta}) \Delta t + \sigma_\theta \Delta Z_t \quad (4)$$

- $\theta_{t+\Delta t}$ and θ_t represent the values of the θ components at times $t+\Delta t$ and t respectively.
- ΔZ_t is the increment of the Wiener process Z_t over the interval Δt , typically simulated as $\sqrt{\Delta t} \times N(0,1)$, adjusted for correlation ρ with ΔW_t .
- The correlation between ΔZ_t and ΔW_t is typically accomplished using Cholesky decomposition.

These equations provide a basis for simulating the paths of μ_t and θ_t over discrete time intervals. By iterating these equations over a sequence of time steps, you can construct simulated trajectories of the processes, which are essential for empirical analysis and application of the model to financial data.

¶ Derivation of the Log-Likelihood Function

To derive the log - likelihood function, we first note the key assumptions:

- The increments $\Delta\mu_t$ and $\Delta\theta_t$ are normally distributed due to the properties of Wiener processes and the central limit theorem for sufficiently small dt .
- We assume the observational data comes in the form of discrete time series, where we observe the outcomes that are thought to be generated by the underlying processes described by our model.

Given these assumptions, the likelihood of observing a specific set of data points, given our model parameters ($\bar{\theta}$, σ_μ , κ , σ_θ and ρ), can be described by the product of the probabilities of each observed increment, under the normal distribution assumption.

Let Y_t represent the observed return at time t , which could be modeled as a function of $\Delta\mu_t$ and $\Delta\theta_t$, or directly related to one of these components. The likelihood function L for observations $\{Y_t\}$ over N time steps given the model parameters is:

$$L(\bar{\theta}, \sigma_\mu, \kappa, \sigma_\theta, \rho \mid \{\Delta R_t\}) = \prod_{t=1}^N f(\Delta R_t; \bar{\theta}, \sigma_\mu, \kappa, \sigma_\theta, \rho) \quad (5)$$

where f is the probability density function of R_t given the parameters.

Assuming normality, and that Y_t directly reflects μ_t and θ_t , f might be expressed in terms of the normal density function with mean and variance derived from the model's parameters and the specific discretization scheme.

The log - likelihood ℓ is the natural logarithm of L , which turns the product into a sum:

$$\ell(\bar{\theta}, \sigma_\mu, \kappa, \sigma_\theta, \rho \mid \{\Delta R_t\}) = \sum_{t=1}^N \log(f(\Delta R_t; \bar{\theta}, \sigma_\mu, \kappa, \sigma_\theta, \rho)) \quad (6)$$

In order to more fully define the log-likelihood and facilitate its estimate, we make the following assumptions:

- Log>Returns Relation:

The observed log>Returns over a time interval Δt are primarily influenced by the long-term drift component μ_t and exhibit mean-reversion as captured by the short-term component θ_t . The total log-return R_t from t to $t+\Delta t$ can then be approximated as a combination of contributions from both factors.

- Expected Log-Return:

The expected log-return over Δt is directly proportional to the drift component, adjusted by a factor that represents the average expected contribution from the mean-reverting component. For simplification, it is assumed this contribution is a function of the difference between the current mean-reverting level and the long-term mean, scaled by the mean-reversion speed κ .

- Variance of Log>Returns:

The variance in log>Returns over Δt incorporates contributions from both the drift and mean-reverting components, adjusted for their correlation. The variance reflects the combined effects of volatilities σ_θ and σ_μ and the correlation ρ between the processes and, for a single time step, can be expressed as:

$$\text{Var}(\Delta R_t) = (\sigma_\mu^2 + \sigma_\theta^2 + 2\rho\sigma_\mu\sigma_\theta)\Delta t \quad (7)$$

- When deriving the log-likelihood function based on this adjusted variance, consider that each observed log-return increment, $\Delta R_t = R_{t+\Delta t} - R_t$ where R_t is the log-return at time t , follows a normal distribution with mean $\mu_t \Delta t$ (assuming the mean return rate for the timestep) and variance

$\text{Var}(\Delta R_t)$ as defined above.

The log-likelihood of observing a series of log-return increments given the model parameters can be expressed as the sum of the log-probabilities of observing each increment under this normal distribution assumption:

$$\ell(\bar{\theta}, \sigma_\mu, \kappa, \sigma_\theta, \rho \mid \{\Delta R_t\}) = -\frac{1}{2} \sum_{t=1}^N \left[\log(2\pi \text{Var}(\Delta R_t)) + \frac{(\Delta R_t - \mu_t \Delta t)^2}{\text{Var}(\Delta R_t)} \right] \quad (8)$$

Where :

- ΔR_t is the observed log-return increment from t to $t+\Delta t$.
- $\mu_t \Delta t$ is the expected log-return increment based on the model.
- $\text{Var}(\Delta R_t)$ is the variance of the log-return increment as derived above.
- The model parameters $(\bar{\theta}, \sigma_\mu, \kappa, \sigma_\theta, \rho)$ influence μ_t , $\text{Var}(\Delta R_t)$ and thus the likelihood of the observed data.

7. References

Adrian, Tobias, and Rosenberg, Joshua. (2008).

Stock returns and volatility: Pricing the short-run and long-run components of market risk.
The Journal of Finance, 63(6), 2997-3030.

Casassus, Jaime, and Collin - Dufresne, Pierre (2005).

Stochastic convenience yield implied from commodity futures and interest rates.
The Journal of Finance, 60 (5), 2283 - 2331.

Longstaff, Francis A., and Schwartz, Eduardo S. (1992).

Interest rate volatility and the term structure: A two-factor general equilibrium model.
The Journal of Finance, 47(4), 1259-1282.

Ribeiro, Diogo R., and Hodges, Stewart D. (2004).

A two-factor model for commodity prices and futures valuation.
EFMA 2004 Basel Meetings Paper.

Schwartz, Eduardo S. (1997).

The stochastic behavior of commodity prices : Implications for valuation and hedging.
The Journal of Finance, 52 (3), 923 - 973.

Schwartz, Eduardo S., and Smith, James E. (2000).

Short - term variations and long - term dynamics in commodity prices.
Management Science, 46 (7), 893 - 911.

Wolfram Language Code

```
(*Dynamics of  $\theta(t)$  and  $\mu(t)$  using the Euler-Maruyama method*)
NextTheta[theta_, thetaBar_, kappa_, sigmaTheta_, deltaZ_, dt_] :=
  theta - kappa * (theta - thetaBar) * dt + sigmaTheta * deltaZ;
NextMu[mu_, theta_, sigmaMu_, deltaW_, dt_] := mu + theta * dt + sigmaMu * deltaW;

LogLikelihoodFunction[params_List,
  observedLogReturns_, dt_, initialTheta_, initialMu_] := Module[
  {theta = initialTheta, mu = initialMu, thetaBar, kappa, sigmaTheta, sigmaMu,
   rho, n = Length[observedLogReturns], logLikelihood = 0, deltaWs, deltaZs,
   totalVariance}, {thetaBar, kappa, sigmaTheta, sigmaMu, rho} = params;
  If[And[sigmaTheta > 0, sigmaMu > 0, -1 < rho < 1],
    (*Generate correlated Wiener process increments*)
    {deltaWs, deltaZs} = Transpose[RandomVariate[MultinormalDistribution[{0, 0},
      {{dt, rho Sqrt[dt] Sqrt[dt]}, {rho Sqrt[dt] Sqrt[dt], dt}}], n - 1]];
    (*Iterate over observed log-returns*)For[t = 1, t < n, t++,
    (*Update theta and mu for the next time step using model dynamics*)
    theta = NextTheta[theta, thetaBar, kappa, sigmaTheta, deltaZs[[t]], dt];
    mu = NextMu[mu, theta, sigmaMu, deltaWs[[t]], dt];
    (*Calculate total variance for the timestep*)
    totalVariance = (sigmaMu^2 + sigmaTheta^2 + 2 rho sigmaMu sigmaTheta) dt;
    (*Accumulate log-likelihood using the direct log-PDF computation*)
    logLikelihood += -0.5 Log[2 Pi totalVariance] -
      ((observedLogReturns[[t + 1]] - mu)^2 / (2 totalVariance));];
  -logLikelihood, (*Return the negative log-likelihood for optimization*)
  Infinity (*Handle invalid parameter values*)];];
```

```
K0 = QuantityMagnitude@
  FinancialData["K0", "Close", {{2010, 1, 1}, {2014, 12, 31}}]
```

```
TimeSeries[]
```

```
K0returns = Prepend[Differences@Log[K0["Values"]], 0];
```

```
dt = 1 / 252;
```

```
observedLogReturns = K0returns;
```

```
(*Adjust the initial conditions based on
  domain knowledge or descriptive analysis of the data*)
initialTheta = Mean[observedLogReturns];
initialMu = observedLogReturns[[1]];
```

```
(*Refine initial parameter guesses based on reasonable expectations*)
initialParamsGuess = {0.001, -0.05, 0.01, 0.02, 0.01}; (*Example adjustments*)
```

```
(*Optimization constraints: Adjust based on model requirements*)
constraints = {0 < sigmaTheta ≤ 1, 0 < sigmaMu ≤ 1, -1 < rho < 1, kappa < 0};
```

```
(*Perform optimization with refined parameters and constraints*)
```

```
optimizedResult =
```

```
  NMinimize[{LogLikelihoodFunction[{thetaBar, kappa, sigmaTheta, sigmaMu, rho},
    observedLogReturns, dt, initialTheta, initialMu], Sequence@@constraints},
    {thetaBar, kappa, sigmaTheta, sigmaMu, rho},
    Method → "DifferentialEvolution"];
```

```
(*Display optimized parameters*)
```

```
optimizedResult
```