Optimal Bayesian Betting & Favorable Games

by

Zhengyu Tang

Department of Statistical Science
Duke University

Defense Date: March 28, 2024

Approved:

David Banks, Supervisor

Simon Mak

David Ye

Thesis submitted in partial fulfillment of the requirements for the degree of
Master of Science in the Department of Statistical Science
in the Graduate School of Duke University
2024

ABSTRACT

Optimal Bayesian Betting & Favorable Games

by

Zhengyu Tang

Department of Statistical Science
Duke University

Defense Date: March 28, 2024

Approved:

David Banks, Supervisor

Simon Mak

David Ye

An abstract of a thesis submitted in partial fulfillment of the requirements for the degree of Master of Science in the Department of Statistical Science in the Graduate School of Duke University
2024

Abstract

Trading has always been considered more of an art than a science. With the rise of quantitative finance, high-frequency trading, and the demoralized equity market, there is more than ever a need to understand why specific strategies make a profit and others do not. This work focuses on the why part and tries to distill down the art component of quantitative strategy development to more of a science discipline. At the same time, it tries to lay down the theoretical groundwork for further research. Bayesian statistics is well suited for such a problem because of the inherited uncertainty quantification and its synergy with typical decision science.

To the author's best knowledge, most current works focus on a particular strategy and fail to realize that a strategy consists of various moving, intertwined, yet crucial components that require sophisticated statistical methods to disentangle and correctly attribute the "effect" to each element. Without first entangling, the analysis can quickly fail to uncover the valid underlying profit driver and get lost in the weeds.

We used the first chapter to introduce the most crucial concept of gambling, Kelly's Criterion. We highlighted its connection with the well-known log utility function of money and the theory of utility maximization in optimal decision-making. Further theories were also developed and extended around the criterion to make it suited to the equity market. The second and third chapters each dive deeply into a particular area of quantitative finance. Even though people treat these two areas separately in practice, they follow the same underlying principle. The second chapter focuses on portfolio optimization. Starting with the classical mean-variance portfolio, we extend it with the Kelly Criterion and prove two things: the latter guarantees positive growth. At the same time, the former does not, and a trade-off exists between the Sharpe ratio and the growth rate. The third chapter leaps to pairs trading through the lens of Bayesian statistics and the generative stochastic model. Such formulation offers insight into the profit generation structure and the more statistically "correct" way to conduct such trades.

Dedication

To Cecila and Chloe. Without their relentless support, this work would never be possible.

Contents

Abstract . . . iv

List of Figures . . . viii

1 Introduction . . . 1

2 Kelly Criterion & Favorable Game . . . 3

2.1 The Simplest Game: Coin Toss . . . 3

2.2 Uneven Payoff . . . 5

2.3 Continuous Approximation . . . 6

2.4 Linear Operator: Adding Risk-Free Rate . . . 10

2.5 Performance Benchmarking & Hypothesis Testing . . . 11

2.5.1 Example: Uncovering Advantageous Game . . . 13

2.6 General Portfolio PnL . . . 14

2.7 Performance Decomposition . . . 15

2.7.1 Example: Trading RSI Signal . . . 17

3 Bayesian Portfolio Optimization . . . 20

3.1 Classical Minimum-Variance Portfolio . . . 20

3.2 Multiple Linear Constraints: Target Return Constraint . . . 22

3.3 Soft Constraints . . . 24

3.4 Connections to Kelly . . . 26

3.5 Multivariate Kelly Portfolio . . . 27

3.6 Kelly Portfolio with Constraints . . . 29

3.7 Example: Kelly Portfolio . . . 30

4 Bayesian Pairs Trading . . . 33

4.1 Classical Pairs Trading . . . 33

4.2 Time-Invariant Equilibrium . . . 35

4.3 The Synthetic Portfolio . . . 37

4.4 Time Varying Equilibrium 39
4.5 Marginal Profitability 41
4.5.1 Leveraged Strategy 41
4.5.2 Unleveraged Strategy 42
4.6 Conditional Profitability 43
4.7 Relationship with OU Process 46
Appendix 48
Bibliography 57

List of Figures

2.1 Kelly Coin Toss Simulation 4

2.2 Kelly Continuous Approximation Using Taylor Expansion 9

2.3 Coin Toss Game Performance 13

2.4 Game Posterior Samples for Strategy 1 13

2.5 Edge Comparison Between Two Games 14

2.6 RSI Strategy Trading Demo 17

2.7 Posterior Samples for the RSI Game. 19

3.1 Seven Sector ETFs Portfolio 31

3.2 Constructed Portfolio Performance, Kelly vs. MinVar 32

4.1 The DLM Estimated ε Series 35

1. Introduction

Wall Street is the largest game [27]. An investor's goal in the stock market overlaps with gamblers trying to beat the Blackjack at casinos [26]. Such games require one to make decisions under various uncertainties and objectives. Its relationship to Bayesian decision analysis and the usage of utility function from economics quickly becomes apparent. Define a utility function, average over the predictive distribution with respect to the future rewards conditional on the latest information, then optimize and explore potential decisions. It is the focus of the work and the central theoretical framework that has been introduced, developed, and extended throughout the article.

The theoretical framework can be practically summarized succinctly. A practice can be broken down into two parts, signal and strategy. The signal defines a certain pattern or relationship, often statistical, of one or more assets and provides predictability in the near future. Ideally, signals should provide and can be thought of as a full posterior predictive distribution of a pattern, $p(x_{t+1:t+h}|D_t)$ at a certain time horizon h while incorporating the most recent information D_t and everything relevant previously. The posterior predictive distribution provides both directional information and all the uncertainty measures. For a simple example, if a distribution with positive expectations $E[X_t|D_t] > 0$ and some defined variances $V[X_t|D_t]$, it, at least, tells us the direction in expectation is heading up. The associated risks or uncertainties are, of course, measured by the variances in this context. The relative association between the mean and variances roughly measures the risk-reward trade-off of the opportunity. Another extreme example is $E[X_t|D_t] = 0$; the risk-reward would always be 0, no matter the variances. In other words, one should not invest when the expectation is zero.

Strategies use the comprehensive information provided by the signals to devise trading decisions. Trading involves three major parts: Entry, Exit, and Sizing. Obviously, one can only come up with proper answers to the three major components when given proper information. And we will soon see these are tightly intertwined with each other; changing the entry would change the respective sizing. For instance, comparing two trading opportunities with symmetric distributions, one with \$1 expectation and \$5 variances, and the other with \$1 expectation but \$10 variances.

One should certainly favor the first one over the second.

Given the above framework, the following sections formalize all concepts working backward. We start with the definition of a favorable game and introduce the famous Kelly Criterion and its application to sizing and performance benchmarking in section 2. We also introduced a probabilistic framework for detecting a favorable game and generalized the analysis of a favorable game to any arbitrary games given a fixed utility. Such theoretical development allows one to conduct analysis, inform decision-making, simulate the expected performance, and perform statistical tests of any arbitrary game. It also highlights the importance of decomposition and deconfounding the contribution from various components of a proper strategy, where the section ends with an example. Section 3 ventures into the portfolio optimization paradigm using the concept presented thus far. We found that optimizing portfolios with an expected utility offers a much better interpretation and overall performance while circumventing many of the classical approaches' limitations. Lastly, the section ends with an example of a sample strategy to highlight its importance. Lastly, section 4 extends the Kelly concept of maximizing expected utility to a famous quantitative strategy, pairs trading, and develops a theoretical understanding of the strategy. It starts with a summarization of the classical strategy. We extend it with a probabilistic model and generalize all concepts of the strategy around the model. It sheds light on an array of insights into the strategy. At the end of the section, the distribution of profitability is derived, so one can understand how the profit is made in pairs trading.

2. Kelly Criterion & Favorable Game

Given an advantageous game, sizing means how much one should bet on the game to maximize a certain goal. The question of how big or small a bankroll, C, to bet has a large say in the eventual performance of an entire portfolio. The reasoning is straightforward. For instance, a game has constant expectations, but the associated variability of the realized payoff varies through time and is known beforehand. One sensible strategy is to bet less when the variability is large and vice versa. The resulting performance should obviously outperform one who bets indifferently. The same concept also extends to the stock market, which is a lot like a casino. For example, At every interval, one is quoted on various instruments and has to make a decision(s) on how much of the bankroll, f, to bet in expecting the odds in one's favor. It sounds fascinating; however, we need to have a proper way to define and measure the odds and various other metrics to spot a winning trade.

2.1 The Simplest Game: Coin Toss

To highlight intuition and core concepts, I am going to use a simple game of coin toss, Z. Say we have an advantageous coin that gives us a rate of landing on heads $1/2 < p < 1$ and an equal payoff $1:1$; we win \$1 on every head and lose \$1 on every tail. Formally, we define our edge as the expectation of the game $E[X]$.

$$Z_t \sim_{iid} \text{Bernoulli}(p)$$

$$X_t = 2Z_t - 1, \quad X_t \in \{-1, 1\}$$

$$E[X_t] = p - (1-p) = p - q, \quad V[X_t] = 4pq$$

The t subscript means different realization of the random variable at different time points. Given the above formulation, we can easily write out the edge and the associated variances. How much should we bet at each play to maximize the long-term return? The Kelly Criterion using the log utility function [17] [3] gives us a set of answers that would lead to various interesting results. Say, we want to maximize the portfolio value $V_t(f)$ with a starting capital V_o at a fixed wager proportion f, or the Kelly fraction, playing the above coin toss game. We have the following derivations.

$$V_T(f) = \prod_{t=1}^{T} V_o(1 + X_t f)$$

$$\ln V_T(f)/V_o = \sum_{t=1}^{T} \ln(1 + X_t f)$$

$$E[\ln V_T(f)/V_o] = g(f) = p\ln(1+f) + q\ln(1-f)$$

$$f^* = \arg\max_f g(f)$$

$$\frac{\partial}{\partial f} g(f) = \frac{p}{1+f} - \frac{q}{1-f} = 0$$

$$f^* = p - q$$

The $g(f)$ function is the expected log-growth rate under a specific wager. Of course, it depends on how much you bet at every game; the log growth rate is different, and the compounded growth rate is going to be even larger in the long run with various betting sizes. The optimal wager f^* is quite intuitive; it says we should bet at $p - q = 0.02$ or 2% of the bankroll every game in our case. On the other hand, for some undesired wagers, the log-growth rate can even be negative under this advantageous game.

Figure 2.1 shows the game's simulation results after $10,000$ times of playing the simple coin

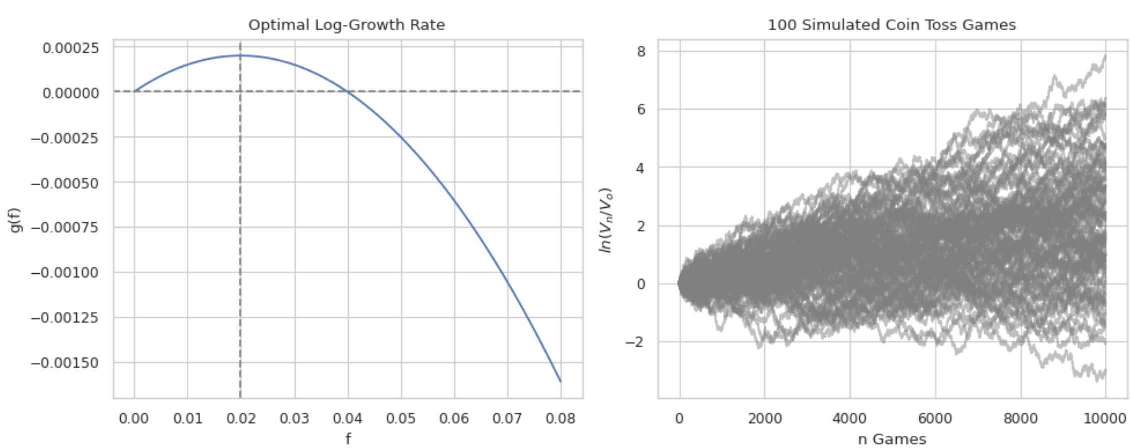

FIGURE 2.1: Kelly Coin Toss Simulation

toss and the relative performance with various Kelly fractions [23]. (Left) We can clearly see the log-growth rate is maximized at 0.02, which is the optimal Kelly fraction. The overall shape is a parabola. Interestingly, it also shows if we bet over 4% or more, the log-growth rate becomes negative. In other words, we would start losing money even if the game is in our favor. (Right) Although we have a small yet clear 2% edge in our little game, the growth of the bankroll is not clear-cut. It follows a stochastic process; it especially follows a discrete Brownian motion with drift. The variances of the bankroll with respect to time scales linearly with time. As the error accumulates over time, one can get extremely lucky or extremely unlucky. However, the mean drift is certainly moving the process upward. We will investigate these properties in later sections.

2.2 Uneven Payoff

Unlike the toy example we introduced in the last section, most games have uneven payoffs. However, the same principle allows. The previous derivation can be easily extended to uneven payout games. Say, instead of the previous even 1 : 1 payouts, we have a d payout ratio that is normalized per unit of loss. Formally, the payout and edge of the new game are defined as:

$$d = \frac{E[X_t|Z_t = 1]}{|E[X_t|Z_t = 0]|}$$

$$\begin{aligned} m = E[X_t] &= E[E[X_t|Z_t]] \\ &= E[X_t|Z_t = 1]p(Z_t = 1) + E[X_t|Z_t = 0]p(Z_t = 0) \\ &= pd - q \end{aligned}$$

We can easily write the equation for the portfolio value $V_t(f)$ using the new definitions.

$$V_T(f) = \prod_{t=1}^{T} V_o(1 + X_t f)$$

$$= \prod_{t=1}^{T} V_o(1 + Z_t f d)(1 - Z_t f)$$

$$\ln V_T(f)/V_o = \sum_{t=1}^{T} \ln(1 + Z_t f d) + \ln(1 - Z_t f)$$

$$E[\ln V_T(f)/V_o] = g(f) = p \ln(1 + fd) + q \ln(1 - f)$$

$$\frac{\partial}{\partial f} g(f) = \frac{pd}{1 + fd} - \frac{q}{1 - f} = 0$$

$$f^* = (pd - q)/d = m/d$$

This is an interesting result: the larger the edge, m, of course, the higher the betting fraction; however, it is reduced or adjusted by the payoff ratios in controlling the overall risk. Expanding further, we can see $f^* = p - q/d$. The larger the payoff ratio, if $d > 1$, the greater the betting fraction. On the other hand, if $d < 1$, we then need p to be much larger than q to make a bet worth a while. No matter what, if the overall edge, m, is negative, no payoffs are worth a while for the bet to happen.

2.3 Continuous Approximation

The simple coin toss game is a great conceptual toy example for examining the Kelly criterion properties. To extend it to the stock market, we have to model the payoffs using continuous distributions because the returns in the actual market are not discrete anymore [29]. Using a continuous distribution offers a lot of advantages in computation. Using the coin toss as an example, assuming we are conducting n tosses every time interval $t \in \{1, \ldots T\}$. As $n \to \infty$, the accumulated results of any interval can be well approximated by a normal distribution. This is the well-known property of the Binomial distribution, which is a sum of a large number of independent Bernoulli random variables. In asymptotic, the Binomial distribution becomes a normal distribution. In fact, the statement comes from the central limit theorem that applies to any distribution with the first two moments defined.

Let's assume a generic game payoff distribution for X_t, where the edge $E[X_t] = m$ and the variance $V[X_t] = s^2$ are defined. Take a time interval t and split it into n equal independent steps,

$$X_t = \lim_{n\to\infty}\sum_{i=1}^{n} X_i, E[X_i] = m/n, V[X_i] = s^2/n$$

, because of the independence. Using this, we write out the usual portfolio value equation.

$$\begin{aligned}
G(f) = \ln\left(V_T(f)/V_o\right) &= \sum_{t=1}^{T} \ln(1+X_t f) \\
&= \sum_{t=1}^{T} \int \ln(1+xf)p(x)dx \\
&= \sum_{t=1}^{T} \lim_{n\to\infty}\sum_{i=1}^{n} \ln(1+fX_i) \\
&= \sum_{t=1}^{T} \lim_{n\to\infty}\sum_{i=1}^{n} \ln(1+0) + \frac{1}{1+0} fX_i - \frac{1}{(1+0)^2}\frac{(fX_i)^2}{2!} + \dots \\
&= \sum_{t=1}^{T} \lim_{n\to\infty}\sum_{i=1}^{n} fX_i - \frac{1}{2}(fX_i)^2 + \dots
\end{aligned}$$

The conversion between the integral and the summation follows the typical Reimann integral approximation. We expanded the original utility function using Taylor expansions to keep the first and second-order terms plus some errors. Let's take a one-time interval and examine its properties under asymptotic expectations. It simplifies to a quite intuitive formula.

$$
\begin{aligned}
E[\sum_{i=1}^{n} fX_i - \frac{1}{2}(fX_i)^2 + \dots] &= \sum_{i=1}^{n} E[fX_i] - \frac{1}{2}E[(fX_i)^2] + E[\dots] \\
&= \sum_{i=1}^{n} \frac{fm}{n} - \frac{1}{2}\left[\frac{f^2 s^2}{n} + \frac{f^2 m^2}{n^2}\right] + E[\dots] \\
&= fm - \frac{1}{2}f^2 s^2 + \varepsilon \\
&= fm - \frac{1}{2}f^2 s^2 + o(n^{-1/2}) \\
\varepsilon &= -\frac{f^2 m^2}{2n} + E[\sum_{i=1}^{n} \dots], \quad \lim_{n\to\infty} \frac{\varepsilon}{n^{-1/2}} \to 0
\end{aligned}
$$

In asymptotic, the expectation of the log-growth rate can be approximated using only the edge m and the variances s^2 and is a function of the Kelly fraction f. Keeping f fixed, the larger the edge, the more the overall log-growth rate increases but is downward adjusted by the associated risk. Such formulation offers great interpretability and flexibility because it applies to any distribution or game that defines the first two moments. In addition, the estimation error bound converges at a speed of $\sqrt{n}$, which is the typical central limit theorem convergence rate.

$$
\begin{aligned}
g_\infty(f) &= \lim_{n\to\infty} E[\sum_{i=1}^{n} fX_i - \frac{1}{2}(fX_i)^2 + \dots] \\
&= fm - \frac{1}{2}f^2 s^2 \\
V_\infty(f) &= \lim_{n\to\infty} \mathrm{Var}[\sum_{i=1}^{n} fX_i - \frac{1}{2}(fX_i)^2 + \dots] \\
&= f^2 s^2 \\
\Delta G_\infty(f) &\sim N(g_\infty(f), V_\infty(f))
\end{aligned}
$$

$$
\begin{aligned}
\frac{\partial}{\partial f} g_\infty(f) &= m - fs^2 = 0 \\
f^* &= \frac{m}{s^2}
\end{aligned}
$$

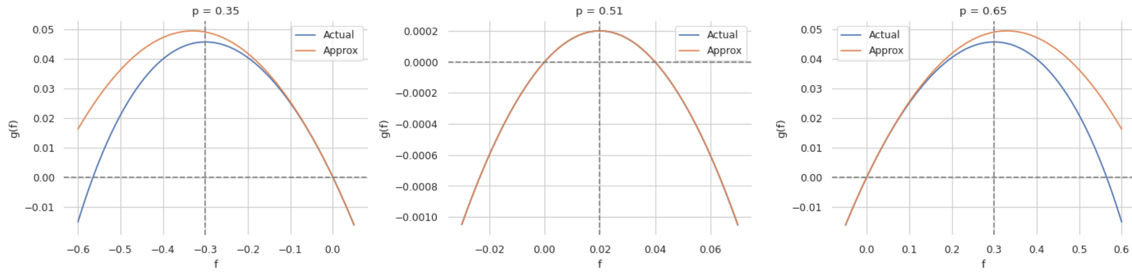

FIGURE 2.2: Kelly Continuous Approximation Using Taylor Expansion

It uses Taylor Expansion and the first two order terms on an advantageous coin-flipping game with a 51% success rate.

Carrying it further, we can write out the full asymptotic stochastic process in terms of Gaussian distribution, with each innovation also following a Gaussian distribution. It shows 3 approximations using the coin toss example in Figure 2.2 with different success probabilities. As we can see, the approximation does quite a good job when the success probabilities center around 50%. On the other hand, when the success probabilities increase or decrease closer to the boundary, the approximation gets more coarse. The optimal fraction especially tends to over-estimate the true value. Nevertheless, the asymptotic formulation provides a nice, systematic way for interpretation. The over-estimation is kind of obvious because the 2nd-order approximation only considers the first 2 moments. In other words, the approximation underestimates the overall risk the closer to the boundary. This can have real-world implications in various scenarios; however, I will not pursue this further in this section. Of course, for robust estimations, one should always consider the full likelihood of X_t, and the typical numerical optimization procedures using the Lagrangian multiplier can be utilized to maximize expected utilities with other constraints.

Now we have arrived at an important and interesting result.

$$G_\infty(f)_t = G_\infty(f)_{t-1} + \Delta G_\infty(f)_t$$

$$\begin{aligned} G_\infty(f)_t | G_\infty(f)_{t-1} &\sim N(G_\infty(f)_{t-1} + g_\infty(f), V_\infty(f)) \\ &\sim N(G_\infty(f)_{t-1} + fm - \frac{1}{2}f^2s^2, f^2s^2) \end{aligned}$$

$$\begin{aligned} G_\infty(f)_t &\sim N(tg_\infty(f), tV_\infty(f)) \\ &\sim N(tfm - \frac{1}{2}tf^2s^2, tf^2s^2) \end{aligned}$$

Each step of the log growth can be approximated asymptotically using a normal distribution with some drift and some variances. It is asymptotic in the sense that we can divide each decision step down to infinitesimal pieces. Note it is valid for any distribution as long as it has the first 2 moments defined. As long as we make a large number of bets, the performance traces will all follow these stochastic equations. This conclusion shouldn't be any of a surprise because of the well-known Central Limit theorem; we just applied it in the time series context. Then, of course, the trace of the log growth $G(f)$ shall follow a random walk with the following conditional and marginal distributions. Figure 2.3 shows an example on how such stochastic trace might pan-out on a simple coin-toss game.

2.4 Linear Operator: Adding Risk-Free Rate

Often times we have multiple choices while gambling or investing. Under a fixed capital, one need to make decision on allocating the capital to different investment vehicles. One of such is interest rate. In theory, one can always put the uninvested capital into a market account and earn the interest rate risk-free. Of course, there are other subtleties in real-world applications, but let's look at how this property affects the Kelly fraction. First, we need to adjust the log-growth equation.

Denote r for the risk-free rate. Then, all derivations follow from the previous section.

$$G(f) = \sum_{t=1}^{T} \ln(1 + (1-f)r + fX_t)$$

$$= \sum_{t=1}^{T} \ln(1 + r + f(X_t - r))$$

$$g_\infty(f) = r + f(m-r) - \frac{1}{2}f^2 s^2, \quad V_\infty(f) = f^2 s^2$$

$$f^* = \frac{m-r}{s^2}$$

Apparently, after adding the risk-free rate, the stochastic process becomes a "mean-reverting" random walk that gravitates back towards the risk-free rate with drift and variances.

There is one pattern we can see here if we treat $Y_t = r + f(X_t - r)$ as one signal random variable or any linear transformation, $Y_t = a + bX_t$. Where X_t is a sum of all instantaneous independent changes of X_i, using the central limit theorem, X_t can well be approximated with a normal distribution in asymptotic as long as n is large. Hence, using the Gaussian properties, any linear transformation of a Gaussian random variable stays Gaussian. We have the following properties.

$$\frac{X_i - m/n}{\sqrt{s^2/n}} \Rightarrow N(0,1)$$

$$g_\infty(f) = E[Y_t] - \frac{1}{2}V[Y_t], \quad V_\infty(f) = V[Y_t]$$

2.5 Performance Benchmarking & Hypothesis Testing

By now, we should realize that the 3 pillars of a strategy, Entry, Exit, and Sizing, are sufficiently summarized by the strategy's edge m, variance s^2, and the proper betting fraction f, at least in asymptotic. Any combination of the tuple (m, s^2, f) forms a different performance. Even if one is playing an advantageous game, varies betting fractions would vary the overall observed performance even to negative. Therefore, when comparing two strategies, we need to consider all these 3 parameters, and the task essentially becomes an inference problem that is widely studied in statistical science.

Take an overly simplified example, assume we have collected performance data for two strategies $j \in \{A, B\}$. The data $D_j = \{(r,f)_{t,j}\}_{t=1}^{T_j}$ can have varying time lengths for different strategies, with each tuple containing a one interval nominal return and the respective invested bankroll. Then, both m and s^2 can be estimated using the Maximum Likelihood Estimation (MLE) by utilizing the asymptotic normality and Markov properties. Let's denote the equity curve to be $V_{0:t}$ for one of the strategies. Then we can write out the full likelihood.

$$\begin{aligned} p(V_{0:T}|f) &= p(V_{1:T}|V_0) = p(V_T|V_{T-1})p(V_{T-1}|V_{T-2})\dots p(V_1|V_0) \\ &= \prod_{t=1}^{T} p(V_t|V_{t-1}) = \prod_{t=1}^{T} N(V_t|V_{t-1} + g_\infty(f), V_\infty(f)) \\ &= \prod_{t=1}^{T} N(V_t|V_{t-1} + fm - \frac{1}{2}f^2s^2, f^2s^2) \end{aligned}$$

$$\begin{aligned} \ell(\theta|V_{0:T}, f) &= \ln[p(V_{0:T})], \quad \theta = \{m, s^2\} \\ &= -\ln(fs\sqrt{2\pi}) - \sum_{t=1}^{T} \frac{1}{2fs}(V_t - V_{t-1} - fm + \frac{1}{2}f^2s^2)^2 \end{aligned}$$

$$\hat{\theta} = \arg\max_{\theta} \ell(\theta|V_{0:T}, f)$$

The estimations follow the frequentist CLT bounds and can thought of as a special case of the AR(1) model by enforcing the first-order parameter equals to 1. $\hat{\theta}_A$ and $\hat{\theta}_B$ can be properly compared using various statistical methods in hypothesis testing.

This framework can be easily extended to time-varying fractions. Even better for robust inference, we can apply Bayesian analysis conditional on known fractions. Then, posterior distributions can be sampled using MCMC using the following model.

$$r_{t,j}|f_{t,j} \sim N(f_{t,j}m_j - \frac{1}{2}f_{t,j}^2 s_j^2, f_{t,j}^2 s_j^2)$$

$$m_j \sim N(m_o, s_m^2), \quad 1/s_j^2 \sim \text{Gamma}(\frac{n_o}{2}, \frac{n_o s_o^2}{2})$$

, where we gave a prior distribution for m_j and s_j^2 individually. The posterior distribution $p(m_j|D_j)$ and $p(s_j^2|D_j)$ can be used to compare among different strategies.

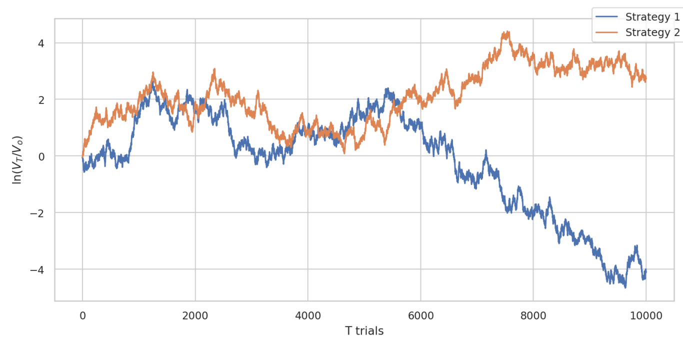

FIGURE 2.3: Coin Toss Game Performance

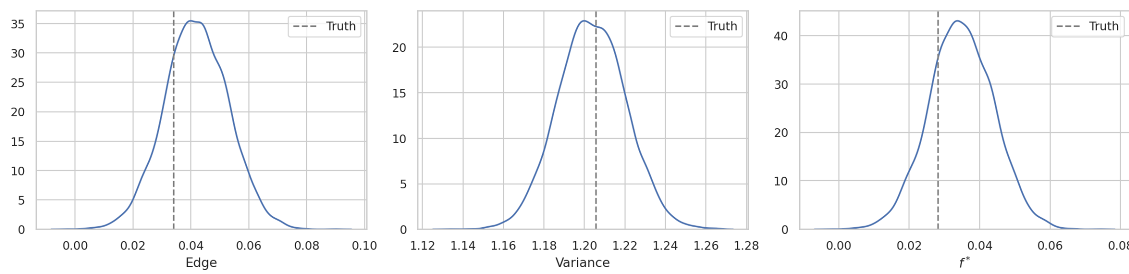

FIGURE 2.4: Game Posterior Samples for Strategy 1

2.5.1 Example: Uncovering Advantageous Game

Let's walk through an interesting example, where we play a coin toss game with a success rate of 47% but a payoff ratio of 1.2 instead. Hence, this game has a positive edge of $E[X] = 0.034$, a variance of $V[X] = 1.206$, and the optimal Kelly fraction $f^* = 2.82\%$. However, we will not bet at the optimal fraction but 4% just to make the problem interesting. Figure 2.3 is the log performance trace after $10,000$ games (Strategy 1). Obviously, it is not performing well; we lost almost all of our bankroll. It is partly due to the fact that we are betting more than the optimal fraction, and part was just pure luck. However, are the odds not on our side, or it is, are we taking too much risk? Apparently, the ground truth says the game is advantageous. In fact, it is. Using the Bayesian model and MCMC, we show the posterior samples in Figure 2.4 along with the ground truth. It matches the expectations and uncovers the true parameters with uncertainty measures, even from a bad performance.

Now, we introduce the same performance trace from a different game, Strategy 2. Its game has

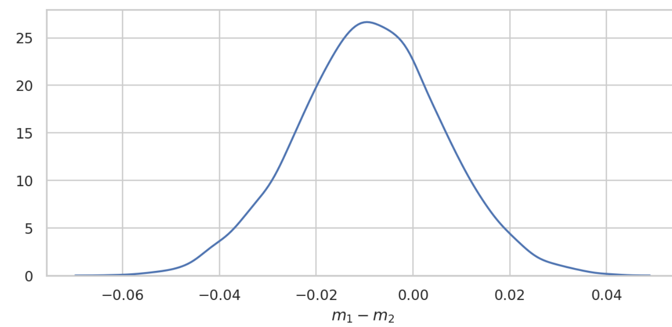

FIGURE 2.5: Edge Comparison Between Two Games

a 51% success rate with an even payoff. Hence, the edge is 0.02, the variance is 1.0, and the optimal fraction is 2%, where we bet at 4%. Obviously, it has a smaller edge, but by chance, strategy 2 produced a better performance curve than Strategy 1. In other words, we should prefer Strategy 2 over Strategy 1 if we have the ground truth. However, we don't have it in practice, but we can utilize the two posterior distributions of the edge $p(m_1 - m_2|D_t)$ to conduct a hypothesis test. It is shown in Figure 2.5. We see there is no statistical significance between the two underlying games. Although, we know the actual difference in edge is $m_1 - m_2 = 0.014$. At least about the $10,000$ observed samples among each strategy's trace, we are not able to tell the difference.

2.6 General Portfolio PnL

In this section, we derive the general formula for calculating a portfolio return (PnL) using log prices that consist of multiple securities. It comes in handy in various simulation and backtesting exercises, and, most importantly, it reveals a simple yet insightful relationship between signal and strategy. Under the usual $t+1$ trading restriction, the positions we enter at t realize at the next time interval, $t+1$, which we will pay special attention to. We use a pair of assets, (X,Y), to highlight the intuition; it, of course, generalizes to multiple assets. We can simply set one of the assets with 0 return and risk to emulate risk-free cash holding.

At time t, we allocate $w_{y,t}$ percentage of total capital to asset Y and $w_{x,t}$ percentage of total capital to asset X, with the constraint $w_{y,t} + w_{x,t} = 1$. And when $w_{\cdot,t}$ is negative, it can be interpreted as shorting the asset, while positive value, of course, means longing for the asset. Using the first principle, we can rewrite the discrete PnL of a portfolio $P_{t+1|t}$ as follows.

$$
\begin{aligned}
\Delta P_{t+1|t} &= q_{y,t}\Delta y_{t+1} + q_{x,t}\Delta x_{t+1} \\
&= (w_{y,t}\frac{C}{y_t m_y})\Delta y_{t+1} + (w_{x,t}\frac{C}{x_t m_x})\Delta x_{t+1} \\
&= C\left[w_{y,t}R_{y,t+1}/m_y + w_{x,t}R_{x,t+1}/m_x\right] \\
C^{-1}\Delta P_{t+1|t} &= \Delta\tilde{P}_{t+1|t} = w_t^{\intercal} R_{t+1}
\end{aligned}
$$

$$
w_t = \begin{bmatrix} w_{y,t} & w_{x,t} \end{bmatrix}^{\intercal} \quad R_{t+1} = \begin{bmatrix} R_{y,t+1}/m_y & R_{x,t+1}/m_x \end{bmatrix}^{\intercal} \quad R_{x,t+1} = \frac{\Delta x_{t+1}}{x_t}
$$

, where $q_{\cdot,t}$ denotes the quantities held in the portfolio for a particular asset at time t, C for total investable capital, $m_{\cdot}$ for the leverage multiplier for a particular asset (i.e., future contract multipliers and 1 for equities), and $R_{y,t+1} = \Delta y_{t+1}/y_t$ is the usual nominal return. The above equation is quite straightforward; it states that the Portfolio PnL is a linear combination of all assets' returns weighted by the respective betting fractions and scaled by the respective invested capital. Note two things: (a) it generalizes two vectors of p-dimension, $\{w, R\}_t \in \mathbb{R}^p$, and (b) we used $\Delta\tilde{P}_{t+1|t}$ to underscore that this is the unit profit in terms of $1 being invested. Lastly, the unit portfolio value at time t forms a trace that follows below.

$$
\tilde{P}_t = \sum_{i=0}^{t-1} w_i^{\intercal} R_{i+1}
$$

The performance obviously consists of a product of two components: bets and realized returns. The investor has only control of the first part, where the decision comes into play, and not the second part.

2.7 Performance Decomposition

One natural follow-up question for the previous section is where or how we come up with w_t at each decision time point. By introducing a function $h(X_t) \Rightarrow \mathbb{R}^p$ that maps some "feature" to a p-dimension vector that matches with the return vector. We call the feature X_t the signal, which can be basically anything. One should separate the analysis between the algorithm and the signal.

There is no use if the signal has no predictability, no matter how one trades with respect to it. On the other hand, how one trades a particular signal greatly affects the overall performance, as we have already seen in section 2.1. Hence, there is no "implicit" signal verification through Backtesting; the overall strategy performance is an amalgam of the two or more. Often, separating the concerns offers a lot of insights on potential improvements.

An ideal signal should provide predictability of future returns. In other words, the signal should be highly correlated with the future returns in both the direction and the magnitudes. So, using this logic, whatever the instrument one trades, one should expect a consistent profit by trading directly proportional to the signal. Mathematical speaking,

$$\Delta\tilde{P}_t = h(\text{Signal}_t)^\intercal R_t$$

$$E[\Delta\tilde{P}_t] > 0$$

, in asymptote, $h(.)$ for some transformation of the signal. One simple example, $h(\text{Signal}_t) = -\ln(\text{Price}_t/\text{Price}_{t-1})$, the negative log return of the previous period. The implicit assumption of this signal says that one believes that the next period's return is negatively correlated with the current period's return. In other words, one can think of $h(.)$ as a mapping between the signal and the position size.

Autocorrelation can affect the analysis above using the expectation directly. $\Delta\tilde{P}_t$ can be highly auto-correlated. Hence, instead of the expectations, we can use the cumulative equity curve, $\tilde{P}_t$, to inspect whether the curve is going up or down. This provides a simple proxy for the signal strength and the baseline performance.

Once the signal is proven to be helpful, we can start thinking about how to trade (sizing, holding length, etc.), various risk management measures (interventions like stop loss, signal filters, etc.), and accounting for executions (slippage, transaction costs, etc.). This component can only be correctly done using a Backtesting system. Most importantly, by adding these additional considerations, one should expect the overall strategy performance to improve over the most naive approach — trading directly proportional to the signal.

To sum it up, one needs to pay more attention to the actual trading signal by analyzing and

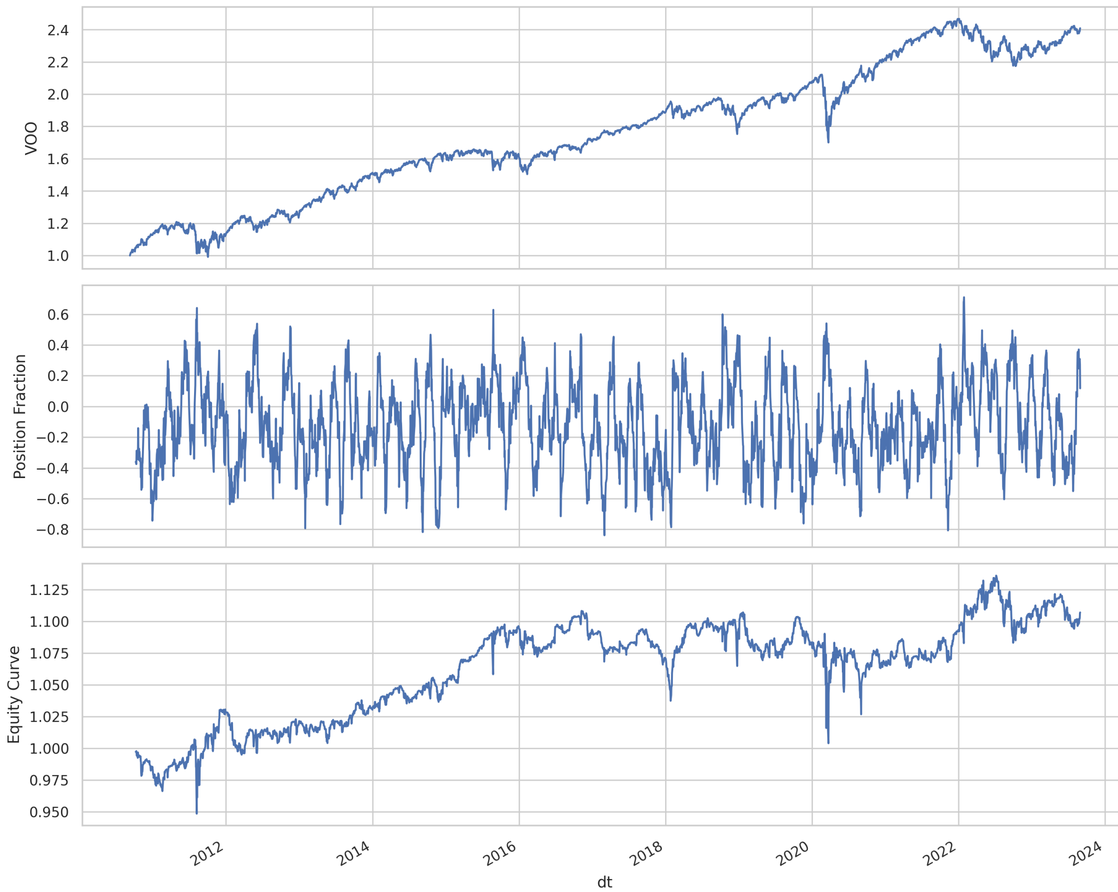

FIGURE 2.6: RSI Strategy Trading Demo

The demo uses VOO, the Vanguard S&P 500 ETF, as an over-simplified example for studying trading signal decomposition. Here the trading signal is the RSI indicator.

understanding the signal through various lenses (how strong is the signal, what is beta exposure concerning the market, and what is the baseline return of the signal, etc.). Too many components contribute to the overall performance and create a permutation of choices; if one doesn't know which components affect the most and their relative potential for improvement, it is easy to get lost in the myriad of options. It leads to a hodgepodge strategy by simply including everything they can think of.

2.7.1 Example: Trading RSI Signal

To highlight the importance of decomposition in the trading strategy analysis, we will walk through a case using the Relative-Strength-Index (RSI) to trade VOO, an ETF, for buying the S&P

500 index as an over-simplied example. We used the year between 2010 and 2023 daily data, calculated the signal, and traded negatively proportional to the signal. Then, the equity curve $\tilde{P}_t$ can be easily calculated using the equations from section 2.6. The RSI indicator is defined as follows, which is defined between $\in [0,1]$, and r_t is the log return of the security. The lookback window, in this case, was chosen as 20 days.

$$\begin{aligned}
\text{RSI}_t &= \frac{\overline{\text{Win}}}{\overline{\text{Win}} + \overline{\text{Loss}}} \\
&= \frac{\frac{1}{T}\sum \max(0, r_t)}{\frac{1}{T}\sum \max(0, r_t) - \frac{1}{T}\sum \min(0, r_t)} \\
&= \frac{\sum \max(0, r_t)}{\sum |r_t|}
\end{aligned}$$

Then, the corresponding position, $h(\text{RSI}_t)$, is a simple transformation to the space of $\in [-1, 1]$. It means we are shorting the asset when the RSI is larger than 0.5 and longing the asset when it is less than 0.5.

$$h(\text{RSI}_t) = -2(\text{RSI}_t - 0.5)$$

The market had a great run during the period, and it is a straightforward case to highlight that the underlying game has a positive expectancy, as will become apparent by investigating the strategy performance shown in Figure 2.6. The market had performed 2.4× over the past 10 years or resulted in an 11% annualized return; on the other hand, the strategy traded using the RSI indicator only resulted in a 10% total return for the entire span, clearly inferior to directly buying the index. The question that often plagues the quantitative investor is: What went wrong? Is the signal not a good indicator for future performance, or is the underlying game not favorable? In this case, the game is clearly favorable as we know the ground truth, and trading a single asset disentangled all the confounders. Furthermore, the posterior samples for the game in Figure 2.7 also reinstated the fact that the mean annualized return is around 10% with an annualized volatility of 16%. Hence, we can conclude the poor performance of the strategy is due to the RSI signal.

Another interesting fact is that using the posterior sample, the optimal Kelly fraction is 4 with $[0.85, 7.72]$ for the 95% CI, which suggests the investor buys the index directly on leverage if using

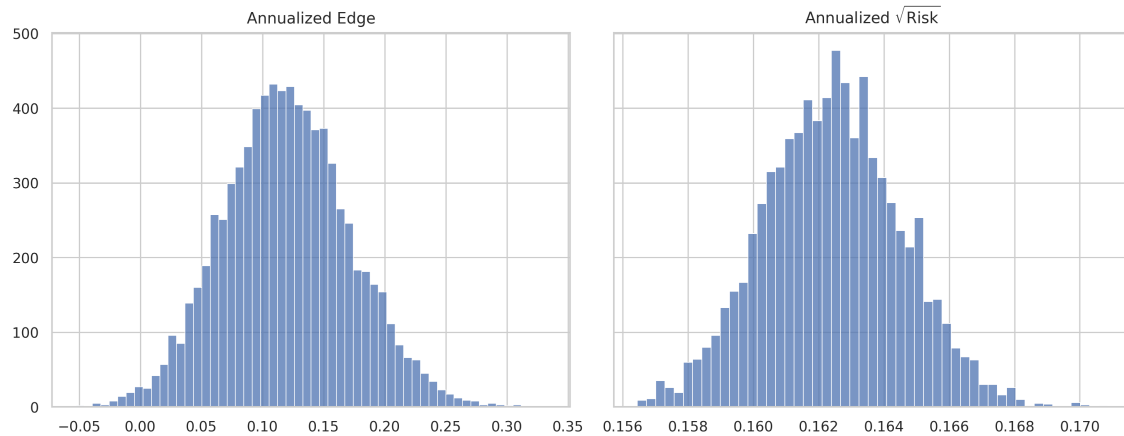

FIGURE 2.7: Posterior Samples for the RSI Game.

the mean. The wide interval gives an additional measure of the underlying uncertainty.

3. Bayesian Portfolio Optimization

Portfolio Optimization has always been a major area in finance and asset management. Since the development of Markowitz's minimum-variance portfolio [21] and the intrinsic meaning of diversification popularized to the public, many variants have come to light. However, the fundamental idea of diversification is never changed; invested in a collection of related assets, one can achieve a certain return with lower risk. The idea is obvious and easily proofed if we only consider the first two moments, and the multivariate Gaussian theory section A proved that, through conditioning, the conditioned variances are always decreasing. More specifically, portfolio optimization is essentially a field blend of decision theory and Bayesian analysis. Instead of doing the optimization using expectation, one should maximize a utility function, $U(f|D_t)$, over a set of posterior predictive distributions on the co-movement of the assets, $p(X_{t+h}|D_t)$, given all the information D_t. The portfolio weight f can be arrived at and explored through the posteriors. On the other hand, the disadvantages of managing a portfolio with the minimum-variance lenses are also apparent. To name a few, it is myopic and doesn't consider the auto-correlated, long-term evolvement of assets; the first two moments might, most likely, underestimate the risk and lead to bad bets, and sometimes, a minimum-variance portfolio can have a negative total return. Nevertheless, the minimum-variance portfolio sets the core idea of modern portfolio theory (MPT) and is worth emphasizing. In later sections, we will extend the concept to a broad scope with a focus on pragmatism. The exercise is highly multi-period, multi-variate, and dynamic.

3.1 Classical Minimum-Variance Portfolio

Let's assume the co-movement of a collection of investable assets $X_t \sim [m, C] \in \mathbb{R}^p$ can be well summarized by the multivariate means, m, and the covariance matrix, C, $C^{-1} = A$ the precision matrix, and they are time-invariant, just to make the problem more tangible for the moment. Then, denote $f_t = [f_1 \dots f_p]^\intercal \in \mathbb{R}^p$ be the invested fractions overall assets at time t. We want to find f_t that minimizes the total portfolio variances with a sum to one constraint. Mathematically, it is shown below.

$$f^* = \min_f \frac{1}{2} f^\intercal C f$$

$$\text{s.t. } f^\intercal \mathbf{1} = 1$$

Here, we removed the subscript t because the mean and covariance don't change through time; so, the optimization becomes a myopic one-period task and is equal for all time periods. It greatly simplifies the problem. Solve the optimization problem using the Lagrangian multiplier; we have the following result.

$$\mathscr{L}(f,\lambda) = \frac{1}{2} f^\intercal C f - \lambda (f^\intercal \mathbf{1} - 1)$$

$$\frac{\partial}{\partial f}\mathscr{L} = f^\intercal C - \lambda \mathbf{1}^\intercal = 0$$

$$f = \lambda C^{-1}\mathbf{1} = \lambda A \mathbf{1}$$

$$\frac{\partial}{\partial \lambda}\mathscr{L} = -f^\intercal \mathbf{1} + 1 = 0$$

$$\lambda = 1/(\mathbf{1}^\intercal A \mathbf{1})$$

We can easily see, the minimum-variance portfolio fractions $f^* = A\mathbf{1}/(\mathbf{1}^\intercal A \mathbf{1})$. Intuitively, the fraction element is the sum of precision in every row adjusted by the return sum of precision, the denominator. In other words, the asset that has lower variances and higher predictability of other assets, either positive or negative, gets higher weight. Another interesting, elegant fact is that the portfolio variances are simply λ.

$$g(f^*) = f^{*\intercal} m$$

$$= m^\intercal \frac{A\mathbf{1}}{\mathbf{1}^\intercal A\mathbf{1}} = \lambda \mathbf{1}^\intercal A m$$

$$V(f^*) = f^{*\intercal} C f^*$$

$$= \lambda^2 \mathbf{1}^\intercal A C A \mathbf{1} = \lambda^2 \mathbf{1}^\intercal A \mathbf{1}$$

$$= \lambda$$

$$S(f^*) = \frac{g(f^*)}{\sqrt{V(f^*)}} = \sqrt{\lambda}\, \mathbf{1}^\intercal A m$$

$g(.)$ is the expected portfolio growth rate, and $S(.)$ is the Sharpe ratio for investing in the minimum-variance portfolio. As we can certainly see from the equation, it is not restrictively positive even if A is s.p.d.

3.2 Multiple Linear Constraints: Target Return Constraint

One way to avoid the undesirable case of negative returns in a minimum-variance portfolio is to redefine a target return constraint. Although a target return m_o might not be known beforehand or, most likely, varies with time — both introduce complexities to the problem — such an approach works well with the existing Lagrangian framework. The typical way to construct it is as follows.

$$f^* = \min_f \frac{1}{2} f^\intercal C f$$

$$\text{s.t. } f^\intercal \mathbf{1} = 1$$

$$f^\intercal m = m_o$$

We simply add an addition constraint as a function of the fractions. The solution is easily followed using two Lagrangian multipliers, λ, γ.

$$\mathscr{L}(f, \lambda, \gamma) = \frac{1}{2} f^\intercal C f - \lambda(f^\intercal \mathbf{1} - 1) - \gamma(f^\intercal m - m_o)$$

$$\frac{\partial}{\partial f}\mathscr{L} = f^\intercal C - \lambda \mathbf{1}^\intercal - m_0 m^\intercal = 0$$

$$f = A(\lambda \mathbf{1} + m_o m) = A\xi$$

Both λ and γ can be solved individually through simultaneous equations and some algebra. However, the insight here is that the solution for f is always a function of the precision matrix and a dot product with the vector ξ that consists of all the linear constraints. This naturally leads to a more generic optimization formulation with multiple linear constraints.

$$f^* = \min_f \frac{1}{2} f^\intercal C f$$

$$\text{s.t. } Bf = b$$

$$B \in \mathbb{R}^{k \times p} \quad b \in \mathbb{R}^k$$

Thus, for k number of linear constraints with its respective constraint value vector, c, $Bf = c$ summarizes all of them in one concise format using matrix notation. For instance, the target return example with sum to one constraint can be easily represented using $B = \begin{bmatrix} \mathbf{1} & m \end{bmatrix}^\intercal$ and $c = \begin{bmatrix} 1 & m_o \end{bmatrix}^\intercal$.

Furthermore, the solution set for the generic version also easily follows. We will use $\lambda \in \mathbb{R}^k$ to lump in all Lagrangian multipliers into a vector.

$$\mathcal{L}(f, \lambda) = \frac{1}{2} f^\intercal C f - \lambda (Bf - b)$$

$$\frac{\partial}{\partial f} \mathcal{L} = f^\intercal C - \lambda^\intercal B = 0$$

$$f = AB^\intercal \lambda = A\xi, \quad \xi = B^\intercal \lambda$$

$$\frac{\partial}{\partial \lambda} \mathcal{L} = -Bf + b = 0$$

$$BAB^\intercal \lambda = b$$

$$\lambda = (BAB^\intercal)^{-1} b$$

Hence, the optimal fraction is $f^* = AB^\intercal (BAB^\intercal)^{-1} b$. More intuitively, we can see the respective portfolio growth and variance also follow an elegant form.

$$\begin{aligned} g(f^*) &= f^{*\intercal}m = \lambda^\intercal BAm \\ &= \xi^\intercal Am \\ V(f^*) &= f^{*\intercal}Cf^* = \lambda^\intercal BACAB^\intercal\lambda \\ &= \lambda^\intercal b \end{aligned}$$

The reason we can obtain an analytical solution for the portfolio is that we are working with a convex loss function, minimizing variance, and a set of linear constraints. It gives great computational savings and a tremendous advantage in real-time quantitative systems. Any variation away from such a form introduces a lot more overhead in the optimization process. The computational burden accumulates the smaller the rebalance interval, of course, where most of the time would be spent inverting the covariance matrix, $C^{-1} = A$, and the respective constraint set $(BAB^\intercal)^{-1}$.

3.3 Soft Constraints

In extending the elegant solution derived in the previous sections, we can introduce some "statistical-friendly" types of constraints or losses. Such formulations are widely used in the Ridge and LASSO regression. In other words, instead of introducing some complex inequality constraints, we can place them in the loss function as a penalty term with a pre-defined radius. Hence, the soft constraints. Concretely, in practice, one often penalizes the target return and the portfolio turnover. The former is because we never know the true return beforehand but only an inking of the expected return; so, having a soft constraint around the expected return makes more sense. In fact, it reduces the jumpiness induced by using the hard constraint. The rebalance constraint is also prevalent because rebalancing a portfolio incurs cost in practice in terms of transactions and slippages. On the other hand, a soft constraint on the portfolio turnover makes sense over a hard constraint. It is obviously not every interval we want to perform a fixed fraction of turnover; it is either equal to or less than the expected fraction. Such a problem can be formulated as follows.

$$\begin{aligned} f^* = \min_f \frac{1}{2}f^\intercal Cf + \frac{1}{2}(f^\intercal m - m_o)^2 + \frac{1}{2}c||f - f_o||_2^2 \\ \text{s.t. } f^\intercal \mathbf{1} = 1 \end{aligned}$$

We have m_o as the usual target return, f_o as last interval portfolio fractions, and c as the penalty term of the ideal radius of the p-dimensional sphere formed between f and f_o vectors in Euclidean space. In other words, we want to minimize the whole equation. It gets larger when the total variance is larger, the return is away from the target return, or when a large turnover is incurred. The analytical solution is derived as follows.

$$\mathscr{L}(f,\lambda) = \frac{1}{2}f^\intercal Cf + \frac{1}{2}(f^\intercal m - m_o)^2 + \frac{1}{2}c(f-f_o)^\intercal(f-f_o) - \lambda(f^\intercal \mathbf{1} - 1)$$

$$\frac{\partial}{\partial f}\mathscr{L} = f^\intercal C + (f^\intercal m - m_o)m^\intercal + c(f-f_o)^\intercal - \lambda \mathbf{1}^\intercal = 0$$

$$[C + mm^\intercal + c\mathbf{I}]f = \lambda \mathbf{1} + m_o m^\intercal + cf_o$$

$$f = [C + mm^\intercal + c\mathbf{I}]^{-1}[\lambda \mathbf{1} + m_o m^\intercal + cf_o]$$

$$= \Omega\xi$$

$$\frac{\partial}{\partial \lambda}\mathscr{L} = -f^\intercal \mathbf{1} + 1 = 0$$

$$\mathbf{1}^\intercal\Omega[\lambda \mathbf{1} + m_o m + cf_o] = 1$$

$$\lambda = (\mathbf{1}^\intercal\Omega\mathbf{1})^{-1}[1 - m_o\mathbf{1}^\intercal\Omega m - c\mathbf{1}^\intercal\Omega f_o]$$

The solution set follows the usual patterns we had in the hard constraint settings, a dot product between the precision matrix and the constraint vector, ξ. In fact, now the precision matrix is further adjusted by the return vector m and the hypersphere's radius c to form a "new" constrained precision matrix Ω. Furthermore, the solution for λ should also seem familiar. If we set m_o and c to zeros, we get back the minimum variance solution. Lastly, the portfolio growth rate and the corresponding variance are in the usual quadratic form.

$$g(f^*) = f^{*\intercal}m = \xi^\intercal\Omega m$$

$$V(f^*) = f^{*\intercal}Cf^* = \xi^\intercal\Omega C\Omega\xi$$

$$S^2(f^*) = m^\intercal\Omega\xi(\xi^\intercal\Omega C\Omega\xi)^{-1}\xi^\intercal\Omega m = m^\intercal\tilde{\Omega}m$$

$S^2(.)$ is the Sharpe ratio squared, where we can see it also follows the usual quadratic form with the precision matrix modified by the corresponding constraints.

3.4 Connections to Kelly

In MPT and Capital Asset Pricing Model (CAPM), we use the Sharpe ratio to determine the risk-reward of the particular investment, where the investments that have the largest Sharpe ratio form a tangent respect to the efficient frontier [21] [24]. For a risk-averse investor and a target return, one should always prefer a portfolio with lower risk that is measured in terms of volatility. In fact, we can see here that the Kelly portfolio is indeed related to this concept but extends MPT to a wider scope. Let's denote the Sharpe ratio as $\tilde{S} = (m-r)/s$, and by plugging back in the optimal Kelly fraction f^* to the expected log-growth function $g_\infty(f)$ using the results derived from section 2.4, we have the following results.

$$
\begin{aligned}
f^* &= \frac{m-r}{s^2} \\
g_\infty(f^*) &= r + f^*(m-r) - \frac{1}{2} f^{*2} s^2 \\
&= r + \frac{(m-r)^2}{s^2} - \frac{1}{2}\frac{(m-r)^2}{s^2} \\
&= r + \frac{(m-r)^2}{2s^2} \\
&= r + \tilde{S}^2/2 \\
V_\infty(f^*) &= f^{*2} s^2 = \frac{(m-r)^2}{s^4} s^2 \\
&= \tilde{S}^2
\end{aligned}
$$

So, the performance of a strategy that follows the optimal Kelly fraction solely depends on the Sharpe ratio in both its expectations and variances. The larger the Sharpe ratio, the larger the log growth rate it should be. In particular, $g_\infty(f^*)$ is the expected log growth rate per time interval, say yearly, which is the sum of the risk-free rate and the Sharpe ratio. Assuming zero interest rate, a 1 Sharpe portfolio would grow the portfolio 50% per year on average. In other words, investing using the Kelly fraction, the Sharpe Ratio, has an interpretation of return implicitly.

3.5 Multivariate Kelly Portfolio

The result we have derived so far, section 2.2, can be easily extended to multivariate cases, which comes in handy in the typical portfolio optimization context. [18], [22], and [23] had explore this subject in various angle. Here, we try to consolidate the results and bring new insights into the Sharpe ratio. Say, we have p choices or assets, the Kelly fraction is then a vector $f = [f_1, \dots, f_p]^\intercal$. The game also becomes a p-dimensional random variable, $X_t \in \mathbb{R}^p$, with a defined mean and covariance or precision matrix, $E[X_t] = m \in \mathbb{R}^p, V[X_t] = C, C^{-1} = A \in \mathbb{R}^{p\times p}$.

$$G(F) = \sum_{t=1}^{T} \ln(1 + r + f^\intercal (X_t - r\mathbf{1}_p))$$

$$g_\infty(f) = r + f^\intercal (m - r\mathbf{1}_p) - \frac{1}{2} f^\intercal C f$$

$$V_\infty(f) = f^\intercal C f$$

$$\begin{aligned} f^* &= \arg\max_f \, g_\infty(f) \\ &= C^{-1}(m - r\mathbf{1}_p) = A(m - r\mathbf{1}_p) \end{aligned}$$

f^* is the usual optimal Kelly fraction that maximizes the asymptotic log growth rate g_∞, and V_∞ is asymptotic variances. By plugging back in the optimal Kelly fraction, we have interesting forms for the growth rate and variances, respectively. Note, here, the $\tilde{S} = \sqrt{\tilde{m}^\intercal A \tilde{m}}$ is the multivariate version of the Sharpe ratio adjusted by the risk-free rate on the game X_t. It is still a scale but is calculated on the entire portfolio.

$$
\begin{aligned}
g_\infty(f^*) &= r + (m - r\mathbf{1}_p)^\intercal C^{-1}(m - r\mathbf{1}_p) - \frac{1}{2}(m - r\mathbf{1}_p)^\intercal C^{-1} C C^{-1}(m - r\mathbf{1}_p) \\
&= r + \frac{1}{2}(m - r\mathbf{1}_p)^\intercal C^{-1}(m - r\mathbf{1}_p) \\
&= r + \frac{1}{2}\tilde{m}^\intercal A \tilde{m} = r + \tilde{S}^2/2 \quad \tilde{m} = m - r\mathbf{1}_p
\end{aligned}
$$

$$
\begin{aligned}
V_\infty(f^*) &= f^{*\intercal} C f^* \\
&= (m - r\mathbf{1}_p)^\intercal ACA(m - r\mathbf{1}_p) \\
&= (m - r\mathbf{1}_p)^\intercal A(m - r\mathbf{1}_p) \\
&= \tilde{m}^\intercal A \tilde{m} = \tilde{S}^2
\end{aligned}
$$

We see that both the log-growth rate and the variances can be written in terms of the precision matrix and a quadratic function between the precision matrix and the respective risk-free-rate-adjusted returns $\tilde{m}$. Note, don't include the risk-free rate in the covariance matrix because it would cause the matrix inversion to be not defined due to the zero variances or singularity.

Nearly identical to the univariate version, the portfolio invested using the optimal Kelly fraction behaves directly proportional to the Sharpe ratio; the larger, the better. It offers another theoretical reason that one should seriously consider the Sharpe ratio in selecting an appropriate portfolio if maximizing growth is the objective, which always seems to be the case for one playing the game to win. Furthermore, we define the portfolio Sharpe ratio $S_\infty(f^*)$ trading at the optimal Kelly fractions.

$$
\begin{aligned}
S_\infty(f^*) &= \frac{g_\infty(f^*) - r}{\sqrt{V_\infty(f^*)}} \\
&= \frac{\tilde{S}^2/2}{\sqrt{\tilde{S}^2}} \\
&= \frac{1}{2}\tilde{S}
\end{aligned}
$$

The portfolio Sharpe states that the trace would produce a ratio that is half of the game's Sharpe ratio. No wonder it is hard to come up with a portfolio with a high Sharpe. Note, $\tilde{S}$ is not a function of f; it is purely Sharpe from the investable securities. In other words, if the underlying assets that are available for investment don't have a high Sharpe, it is even harder for the portfolio to have a high Sharpe, or the linear combination and the relationships through A need to compensate for a larger reduction in variances to make it worth a while. It certainly makes sense intuitively, but it also contradicts the wildly accepted diversification concept. A maximizing growth rate strategy is not necessarily a maximizing Sharpe strategy that the literature adores.

Further development of the portfolio Sharpe ratio $S_\infty(kf^*), k > 0$, by taking on a constant scaled Kelly fraction, offers additional insights between the Sharpe ratio and the portfolio growth rate. The relationship is linear and is highlighted below.

$$S_\infty(kf^*) = (1 - \frac{1}{2}k)\tilde{S}$$

, where we see one is essentially making trade-offs between maximizing portfolio growth rate by taking the optimal Kelly fraction or maximizing the Sharpe ratio instead.

Once again, the results we came up with using a 2nd-order Taylor approximation of the expected log-utility function become coarse and unstable the larger the dimensionality but, nevertheless, offer an insightful picture. And the convex structure offers a great advantage in optimization. Hence, either one can perform the approximation in a finer interval or should use the full likelihood to estimate the parameters in practice.

3.6 Kelly Portfolio with Constraints

Further extending the Kelly portfolio to include the portfolio turnover penalty and writing the entire problem in terms of the optimization framework, we have the following form.

$$\begin{aligned} f^* &= \arg\max_f \, g_\infty(f) - \frac{1}{2}c||f - f_o||_2^2 \\ &= \arg\max_f \, r + f^\intercal(m - r\mathbf{1}) - \frac{1}{2}f^\intercal C f - \frac{1}{2}c||f - f_o||_2^2 \end{aligned}$$

Note we are using the CLT approximated Kelly to ensure convexity, and instead of minimizing, we are maximizing the penalized utility function, the log growth rate, with portfolio turnover as the penalty. f_o is the last interval portfolio fractions, and c is the hypersphere radius as a soft constraint. From a different perspective, the approximated Kelly portfolio essentially maximizes the total return adjusted by the risk, $f^\intercal C f$. With such a formulation, we can opt out of the sum-to-one constraint; we should leverage when necessary and favorable and invest elsewhere when not favorable. There is no need to confine the total fraction to one, which doesn't make practical sense. Because of the involvement of r, the risk-free rate, serving as the "mean-reverting" baseline, the resulting portfolio fractions would be much more stable. Such treatment reduces the inter-period jumpiness in the usual minimum-variance portfolio because one doesn't need to over-leverage to achieve a certain return. The optimal Kelly fraction is easily derived below.

$$\begin{aligned} \frac{\partial}{\partial f}\mathscr{L} &= m^\intercal - r\mathbf{1}^\intercal - f^\intercal C - c(f - f_o)^\intercal = 0 \\ -f^\intercal(C + c\mathbf{I}) &= -cf_o^\intercal - m^\intercal + r\mathbf{1}^\intercal \\ f^* &= (C + c\mathbf{I})^{-1}(m - r\mathbf{1} + cf_o) \\ &= \Omega\mu \end{aligned}$$

We can see the optimal Kelly fractions also have a great interpretability compared to the minimum-variance portfolio. Ω is, as usual, the constraint-adjusted precision matrix. Invertibility is ensured due to the diagonal additions. And, instead of using ξ, Kelly fractions use μ as a proxy for mean returns. The former highlights the hard constraint boundaries, but the latter underscores the expected returns. The fractions are first adjusted by the risk-free rate, r, and then a portion from the previous interval's fractions, cf_o, is added. The larger the fraction in f_o, the larger the proxy returns act as a mechanism of limiting portfolio turnover.

3.7 Example: Kelly Portfolio

We showcase a simple portfolio construction example using the previously derived Kelly portfolio equations to fully complete the portfolio optimization section. Figure 3.1 plotted all ETFs' respective performance for the past 14 years and the correlation between them. Each of the ETFs tracks a particular industry sector. For example, VCR focuses on consumer discretion, VGT focuses on technology, and VDE focuses on energy. Despite all of them ending at different perfor-

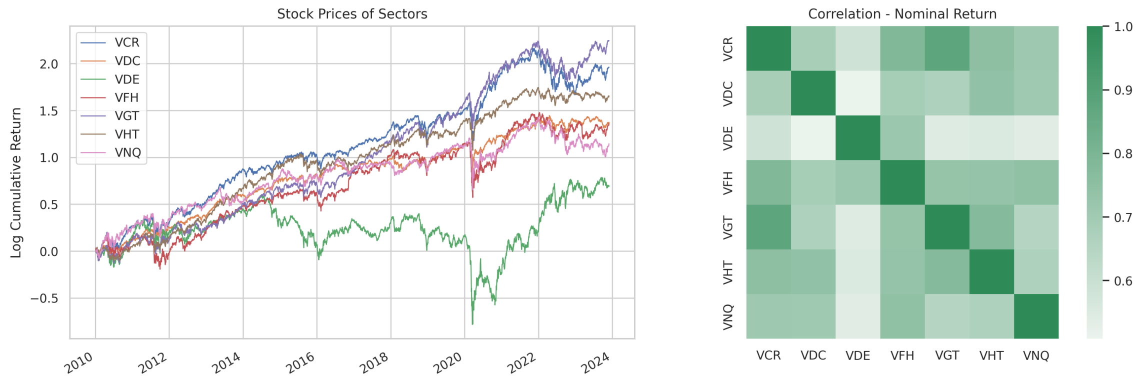

FIGURE 3.1: Seven Sector ETFs Portfolio

The portfolio is constructed using the seven sector ETFs in this example. (Left) Log cumulative curves are plotted for each ETF with different colors. (Right) The covariance matrix calculated the nominal returns.

mances at the beginning of 2024, we can certainly discern many co-movements along the way. The 2020 COVID period is one of the most noticeable; all of them experienced a large drop.

The corresponding covariance or precision matrix is easily calculated, given the price series. Using the equation from section 3.5, $f^*_{\text{kelly}} = A(m-r)$, and assuming the annualized risk-free rate to be 2%, the resulting optimal Kelly fractions are given below.

$$f^*_{\text{kelly}} = \begin{bmatrix} 0.80 & 1.52 & -0.39 & -0.52 & 1.18 & 0.75 & -0.74 \end{bmatrix}^\intercal$$

In contrast, the minimum-variance fraction is defined as $f^*_{\text{minvar}} = \lambda A\mathbf{1}$ and has the values of

$$f^*_{\text{minvar}} = \begin{bmatrix} 0.11 & 0.92 & 0.02 & -0.16 & -0.15 & 0.29 & -0.05 \end{bmatrix}^\intercal$$

, where we see f^*_{minvar} tends to have more extreme values compare to f^*_{kelly}. The reason is because of the sparsity induced through the hard constraint.

We can easily calculate the respective unit portfolio performances using f^*_{minvar} and f^*_{kelly} using equations from section 2.6, and make a comparison between the two. This result is shown in Figure 3.2. It is obvious that the Kelly portfolio outperforms every Minimum Variance portfolio and any other ETF component. It uses leverage and balances risk appropriately with good long-term growth. On the other hand, the minimum variance portfolio doesn't balance between growth and the corresponding risk; thus, the sum-to-one hard constraint. Furthermore, The right chart of Figure 3.2 demonstrates the elegance of the asymptotic theory. As we can see, its portfolio performance is

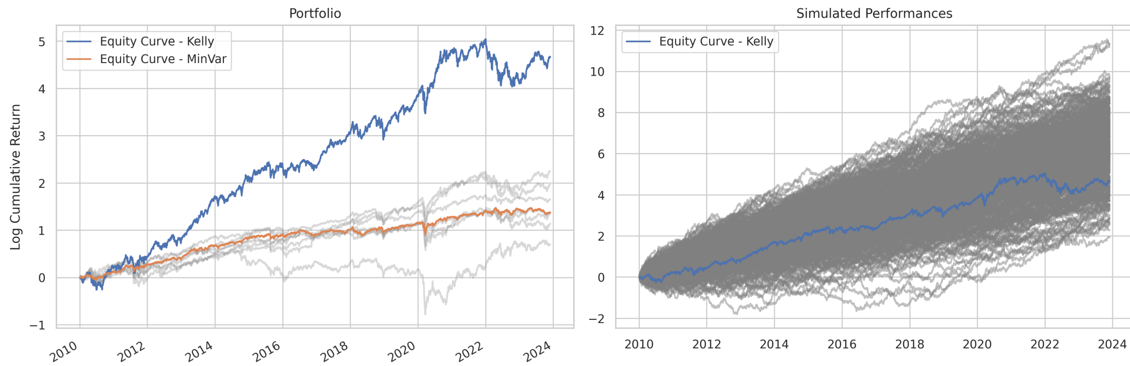

FIGURE 3.2: Constructed Portfolio Performance, Kelly vs. MinVar

Comparison between the Kelly and Minimum Variance methods. (Left) The respective equity curve for each method and each component ETF is in light grey for comparison purposes. (Right) It compares the resulting Kelly portfolio performance, in blue, with the asymptotic theoretical performances, in dark grey, using $g_\infty(f^*)$ and $V_\infty(f^*)$.

well within the expected theoretical boundary. It can be used as a simple hypothesis-testing tool in monitoring model performance. In practice, if the curve runs outside the boundaries, we shall start investigating what causes the deterioration.

Of course, in practice, there are many intricacies to carry out the theory exactly. Namely, when we calculate the covariance matrix, we don't have future data as we used in this example. The covariances matrix has to be derived using historical data and dramatically vary through time. Such an operation introduces a large number of estimation errors to the optimization problem. The optimization only works with the "correct" covariance matrix; if one either under or estimates the true covariances, the resulting fractions will be spurious. The time-varying dimension also introduces difficulties in the estimation task. If one doesn't pay too much attention to how fractions change with time, one will obtain volatile changes that can be practically infeasible. A simple method is to use the portfolio turnover constraint or introduce a penalized term on large fractions. Estimating it is a deep research topic and will not be covered in this section. Interested readers can refer these [9] [16] [6] [31] [14] [5] for further explorations.

4. Bayesian Pairs Trading

Pairs Trading is a trading or investment strategy used to exploit financial markets that are out of equilibrium. According to [12], the concept of pairs trading is relatively simple and can be decomposed into 2 steps. First, find two securities whose prices have moved together in a historical period. Second, the spread constructed from the price between them in a subsequent trading period is monitored. [20] explained the philosophy of Goldman Sachs Asset Management as one of assuming that while markets may not be in equilibrium, over time, they move to a rational equilibrium. Profits come from such a mean-reverting process, and traders are interested in finding or constructing a stable reverting process so as to make consistent profits. They form a trading strategy consisting of a long position in one security and a short position in another security in a predetermined ratio [1]. [19] summarized the work of pairs trading and classified the methods into five categories: Distance approach, Cointegration approach, Time series approach, Stochastic control approach, and other approaches. Nevertheless, this section's focus centers on applying Bayesian decision-making to a pairs trading framework. We start with the classical formulation and then quickly extend it using a probabilistic formulation, allowing a much richer simulation and modeling venue. The strategy's theoretical performance is analyzed both marginally and conditionally. Lastly, we tie the decision-making process with Kelly's criterion of finding a favorable game.

4.1 Classical Pairs Trading

The most widely known pairs trading strategy is conducted on a pair of two assets, which we summarize in this section. For any time point t, we denote a pair of asset prices $\{X,Y\}_t$ indexed by time t, where we assume that the two log price series, $\{\log X, \log Y\}_t$, have a local instantaneous linear relationship with time-invariant parameters $\{\alpha, \beta, \sigma_\varepsilon^2\}$. We also assume ε_t follows a mean-reverting process around zero with parameters $\{k, d\sigma_\varepsilon^2\}$. Thus, the following mathematical relationship,

$$\log Y_t = \alpha + \beta \log X_t + \varepsilon_t \quad \varepsilon_t \sim N(0, \sigma_\varepsilon^2)$$
$$d\varepsilon_t = -k\varepsilon_t dt + d\sigma_\varepsilon dW_t \quad dW_t \sim N(0,1), 0 < k < 1$$

$$d\log Y_t = \beta d\log X_t + d\varepsilon_t$$

, where k is the mean-reverting rate, larger the faster, and $d\sigma_\varepsilon^2$ is the instantaneous volatility of the series per dt. On the other hand, σ_ε^2 is the asymptotic volatility of the series where we have the relationship between the instantaneous. It ensures the series doesn't blow up and preserves the mean-reverting property; in other words, ε_t is an asymptotic bounded process. An example can be seen in Figure 4.1.

$$\sigma_\varepsilon^2 = \lim_{h \to \infty} \text{Cov}(\varepsilon_{t+h}, \varepsilon_t)$$
$$= \frac{d\sigma_\varepsilon^2}{2k}$$

Pairs trading tries to profit from the mean-reverting property of ε_t. Although, marginally, $E[\varepsilon_t] = 0$, but conditionally $E[\varepsilon_{t+h}|\varepsilon_t] \propto \varepsilon_t \exp(-kh)$, the series has a tendency to move up or down proportional to the mean-reverting factor k. The larger the deviation from the origin, the larger the "gravitational force". Hence, the coined term synthetic portfolio as one is trading ε_t directly through hedging between $\{X, Y\}$ in a particular way.

$$d\varepsilon_t = d\log Y_t - \beta d\log X_t$$
$$\frac{1}{1+\beta} d\varepsilon_t = \frac{1}{1+\beta} d\log Y_t - \frac{\beta}{1+\beta} d\log X_t$$
$$d\tilde{\varepsilon}_t = w_{y,t} d\log Y_t + w_{x,t} d\log X_t$$

$$w_t = \begin{bmatrix} \frac{1}{1+\beta} & -\frac{\beta}{1+\beta} \end{bmatrix}^\mathsf{T}$$
$$= \begin{bmatrix} w_{y,t} & w_{x,t} \end{bmatrix}^\mathsf{T}$$

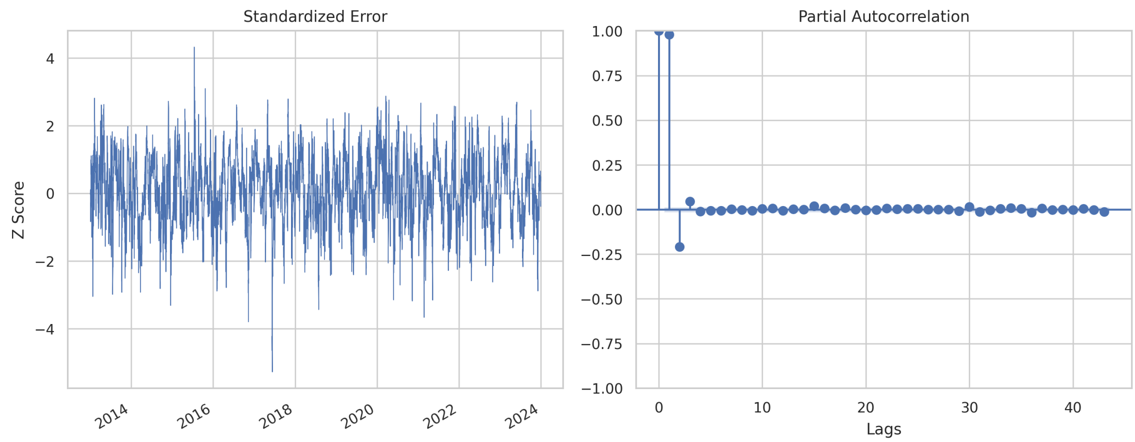

FIGURE 4.1: The DLM Estimated ε Series

The estimated series used the QQQ:SPY pair by the DLM model with an 80-day half-life discount rate. (Left) The standardized error series highlights the mean-reverting behavior by removing the time-varying volatility. (Right) The partial autocorrelation we see in the first order is close to 1 but not equal.

$\tilde{\varepsilon}_t$ is the unit synthetic portfolio as one is investing \$1 into it, where as the original ε_t assumes one invest $\$(1+\beta)$ instead, which greatly simplify things later on. From now on, when we discuss the synthetic portfolio, we always assume it is in \$1 unit. And w_t is known as the hedge ratio because $E[d\tilde{\varepsilon}_t] = 0$ as one hedges out the first-order risk from asset $[X,Y]$. Of course, one can also interpret it as the fraction of capital to be invested in either asset to be invested in the synthetic portfolio.

Lastly, trading-wise, one should short $\tilde{\varepsilon}_t$ when $E[\tilde{\varepsilon}_{t+1} - \tilde{\varepsilon}_t | \tilde{\varepsilon}_t > 0] < 0$ to profit from the mean-reverting property, and vice versa, $E[\tilde{\varepsilon}_{t+1} - \tilde{\varepsilon}_t | \tilde{\varepsilon}_t < 0] > 0$. Ideally, one trades in the same direction as the expectations, and the larger the betting size, the larger the expectations. Of course, there are numerous intricacies in this setup. The estimation of ε_t for one and the time-varying nature of $\{\alpha, \beta, \sigma_\varepsilon^2\}_t$ greatly affect the estimation of other components and the eventually realized profit. We will investigate these in the following sections.

4.2 Time-Invariant Equilibrium

To gain additional insights into Pairs Trading's profit structure, we will impose a general statistical model and work under its assumptions. It extends the classical stochastic formulation and uses a probabilistic framework. First, let us assume the parameters do not vary with time for simplicity. Assuming we know there is an equilibrium μ between two correlated assets $Z = (X,Y)$, we can write out the following Markov process assuming multivariate-normal (MVN) evolution densities.

$$z_t | z_{t-1}, \theta \sim \text{MVN}(\mu + \phi(z_{t-1} - \mu), \Sigma) \quad |\phi| < 1$$
$$z_t = \left[y_t, x_t\right]^\intercal \quad \Sigma = \begin{bmatrix} \sigma_y^2 & \rho\sigma_x\sigma_y \\ \rho\sigma_x\sigma_y & \sigma_x^2 \end{bmatrix}$$

, ϕ is the equilibrium-reverting rate back to μ (smaller the faster), $\sigma_\cdot^2$ is the one-step evolution variances for a particular asset, and ρ is the correlation factor $\in [-1,1]$ between (X,Y). We will group all parameters to keep the bookkeeping easier, $\theta = \{\mu, \phi, \Sigma\}$.

We should notice that z_t follows the usual AR(1) process. By taking the limit on t, we can derive the following result.

$$
\begin{aligned}
z_t &= \mu + \phi(z_{t+1} - \mu) + \varepsilon_t, \quad \varepsilon_t \sim \text{MVN}(0, \Sigma) \\
&= (1-\phi)\mu + \phi z_{t-1} + \varepsilon_t \\
&= (1-\phi^2)\mu + \phi^2 z_{t-2} + \phi\varepsilon_{t-1} + \varepsilon_t \\
&\vdots \\
&= (1-\phi^h)\mu + \phi^h z_{t-h} + \eta_h \\
\eta_h &= \phi^{h-1}\varepsilon_{t-h} + \phi^{h-2}\varepsilon_{t-h+1} + \cdots \phi\varepsilon_{t-1} + \varepsilon_t \\
\lim_{h\to\infty}(1-\phi^h)\mu &= \mu \\
\lim_{h\to\infty} E[\eta_h] &= 0 \\
\lim_{h\to\infty} \text{Var}[\eta_h] &= (\phi^{2(h-1)} + \phi^{2(h-2)} + \cdots + \phi^2 + 1)\Sigma \\
&= \frac{1}{1-\phi^2}\Sigma
\end{aligned}
$$

$$
z_t|\theta \sim \text{MVN}(\mu, \frac{1}{1-\phi^2}\Sigma), \quad |\phi| \leq 1
$$

Therefore, the marginal process also follows a multivariate-normal mean-reverting process with a mean equal to the equilibrium μ and an inflated, yet time-invariant, variance matrix by ϕ with the elliptical shape determined by ρ. Two key takeaways are: (1) ϕ determines how far and how long the process would wander out, and (2) ρ determines how narrow the wondering boundary will be.

4.3 The Synthetic Portfolio

Going back to Pairs Trading, we can derive the hedge ratio, once again, using the multivariate model. Then, the synthetic portfolio ε_t follows.

$$
\beta = \rho\frac{\sigma_y}{\sigma_x}
$$

$$
\varepsilon_t|\theta = y_t - (1-\beta)\mu - \beta x_t
$$

Under the multivariate-normal theory from section A, we know that y_t and x_t individually follow an AR(1) process reading off directly from the diagonal. Therefore, the marginal distribution of the synthetic portfolio $\varepsilon_t|\theta$ has also to be Gaussian. Note that it is only true conditionally on all parameters θ.

$$y_t|\theta \sim N(\mu, \frac{1}{1-\phi^2}\sigma_y^2)$$

$$x_t|\theta \sim N(\mu, \frac{1}{1-\phi^2}\sigma_x^2)$$

$$\begin{aligned}
E[\varepsilon_t|\theta] &= E[y_t - (1-\beta)\mu - \beta x_t|\theta] \\
&= E[y_t|\theta] - (1-\beta)\mu - \beta E[x_t|\theta] \\
&= \mu - (1-\beta)\mu - \beta\mu \\
&= 0
\end{aligned}$$

$$\begin{aligned}
\text{Var}[\varepsilon_t|\theta] &= \text{Cov}[y_t - (1-\beta)\mu - \beta x_t, y_t - (1-\beta)\mu - \beta x_t|\theta] \\
&= \text{Cov}[y_t - \beta x_t, y_t - \beta x_t|\theta] \\
&= \text{Var}(y_t) - 2\beta\text{Cov}[y_t, x_t] + \beta^2\text{Var}[x_t] \\
&= \frac{1}{1-\phi^2}\left[\sigma_y^2 - 2\beta\rho\sigma_x\sigma_y + \beta^2\sigma_x^2\right] \\
&= \frac{1}{1-\phi^2}\left[\sigma_y^2 - 2\rho^2\sigma_y^2 + \rho^2\sigma_y^2\right] \\
&= \frac{\sigma_y^2}{1-\phi^2}(1-\rho^2) \\
&= \sigma_\varepsilon^2
\end{aligned}$$

$$\varepsilon_t|\theta \sim N(0, \sigma_\varepsilon^2)$$

Hence, clearly, we can see the synthetic portfolio ε_t that we are trading marginally follows a variance-bounded stochastic process with a mean 0 and a time-invariant variance σ_ε^2. Especially,

σ_ε^2 has a particular structure; it is damped by the correlation factor ρ, inflated by the convergence rate ϕ, and a multiplier to the dependent asset's volatility σ_y^2. Such a relationship will be useful in selecting the pairs.

On the other hand, it is also important to understand the transitional density of the synthetic portfolio $\varepsilon_{t+1}|\varepsilon_t,\theta$, because that is where the profit is made. To do this, we will utilize the conditional multivariate-normal theorem to derive the exact form. The conditional highlights the mean-reverting nature of the process.

$$\begin{bmatrix}\varepsilon_{t+1}\\ \varepsilon_t\end{bmatrix} \sim \text{MVN}(0, V = \begin{bmatrix}\sigma_\varepsilon^2 & \text{Cov}(\varepsilon_{t+1},\varepsilon_t)\\ \text{Cov}(\varepsilon_{t+1},\varepsilon_t) & \sigma_\varepsilon^2\end{bmatrix})$$

$$\begin{aligned}
\text{Cov}(\varepsilon_{t+1},\varepsilon_t) &= \text{Cov}(y_{t+1}-\beta x_{t+1}, y_t-\beta x_t)\\
&= \text{Cov}(\phi y_t+\varepsilon_{y,t}-\beta\phi x_t+\beta\varepsilon_{x,t}, y_t-\beta x_t)\\
&= \phi\text{Var}(y_t)-2\phi\beta\text{Cov}(y_t,x_t)+\phi\beta^2\text{Var}(x_t)\\
&= \frac{\phi}{1-\phi^2}\left[\sigma_y^2-2\beta\rho\sigma_x\sigma_y+\beta^2\sigma_x^2\right]\\
&= \frac{\phi\sigma_y^2}{1-\phi^2}(1-\rho^2)\\
&= \phi\sigma_\varepsilon^2
\end{aligned}$$

$$\therefore \quad V = \sigma_\varepsilon^2\begin{bmatrix}1 & \phi\\ \phi & 1\end{bmatrix}$$

$$\varepsilon_{t+1}|\varepsilon_t,\theta = N(\phi\varepsilon_t, \sigma_y^2(1-\rho^2))$$

4.4 Time Varying Equilibrium

What would happen if the equilibrium μ_t moves slowly throughout time, say like a random walk? Both $z_t = (y_t, x_t)$ are still mean-reverting around the μ_t, but marginally, none of the series are "mean-reverting" anymore. Say, we have a random walk equilibrium with a mean of 0; we have the following derivations.

$$\mu_t|\mu_{t-1} \sim N(\mu_{t-1}, \sigma_\mu^2)$$

$$\mu_t \sim N(0, t\sigma_\mu^2)$$

The conditional z_t is still mean-reverting but becomes non-stationary if we marginalize over μ_t. The results follow a typical multivariate conditional normal theory.

$$z_t|\mu_t, \phi, \Sigma \sim \text{MVN}(\mu_t, \frac{1}{1-\phi^2}\Sigma), \quad |\phi| \leq 1$$

$$z_t|\phi, \Sigma \sim \text{MVN}(0, \tilde{\Sigma})$$

$$\tilde{\Sigma} = \frac{1}{1-\phi^2}\Sigma + t\sigma_\mu^2$$

We see the covariance matrix is no longer time-invariant and is overly damped by the non-stationary term. The t component comes directly from the random walk equilibrium. However, because ϕ often is really close to 1, the time adjustment is relatively small. Also, we notice the marginal z_t is 0 centered. This shouldn't be any surprise; however, we already know that if a distribution has an expectation of 0, the edge of playing this game is also 0. This observation is central to the Efficient Market Theorem. On the other hand, $\varepsilon_{t+1}|\varepsilon_t, \theta$ doesn't have a 0 means. So it should be apparent by now traders should not be trading using the marginal distribution but the conditional distribution. Once in a while, the time series will wonder to the point that it provides an opportunity, and it always will. The dependency of the trial is where the profit is being made. The inference can be made in real-time using the DLM model introduced in appendix B.

Using the marginal distribution, we can define a notion called the Global beta, $\tilde{\beta}$. It is the $\tilde{\beta}$ if one is estimated using the following linear relationship on the price series.

$$
\begin{aligned}
y_t &= \alpha_t + \tilde{\beta}_t x_t + \varepsilon_t \\
y_t | x_t, \theta_t &\sim N(\alpha_t + \tilde{\beta}_t x_t, \sigma_y^2(1-\rho^2)) \\
\tilde{\beta}_t &= \frac{(1-\phi^2)^{-1}\rho\sigma_x\sigma_y - (1-\phi)^2 t\sigma_\mu^2}{(1-\phi^2)^{-1}\sigma_x^2 - (1-\phi)^2 t\sigma_\mu^2} \\
&= \frac{\beta_t(1-\phi^2)^{-1}\sigma_x^2 - (1-\phi)^2 t\sigma_\mu^2}{(1-\phi^2)^{-1}\sigma_x^2 - (1-\phi)^2 t\sigma_\mu^2} \\
&= \frac{\beta_t a - b}{a - b} \\
\frac{\partial}{\partial b}\tilde{\beta}_t &= \frac{a}{(a-b)^2}(\beta_t - 1)
\end{aligned}
$$

We can think of t as the sample period that is used to run this linear regression. As $t \to 0$, $\tilde{\beta}_t \Rightarrow \beta_t$ converges to the instantaneous beta or the local beta. On the other hand, if $t \to \infty$, the global beta explores. Furthermore, referring to the gradient, due to the time adjustment factor, if $\beta_t > 1$, the global beta $\tilde{\beta}_t$ will be even larger because $\frac{a}{(a-b)^2} > 0$, and vice versa. Another way to interpret this relationship is that the global beta has a larger swing than the local beta.

4.5 Marginal Profitability

4.5.1 Leveraged Strategy

Using the derived distribution for the synthetic portfolio $\varepsilon_t|\theta$, we can show that Pairs Trading is overall a short volatility strategy. Let's assume we have the following naive strategy (for simpler math): At the time t, we enter the position q_t directly proportional to the z-score $z_t = \varepsilon_t/\hat{\sigma}_\varepsilon$ with an opposite sign. In other words, if $z_t = 2$, we short 2x of the synthetic portfolio. Hence, we have

$K_{t+1,t}$ the marginal strategy profit of any bivariate as follows.

$$q_t = -\frac{\varepsilon_t}{\sigma_\varepsilon} = -Z_t$$

$$\begin{aligned}
K_{t+1,t} &= q_t(\varepsilon_{t+1} - \varepsilon_t) \\
&= -Z_t(\sigma_\varepsilon Z_{t+1} - \sigma_\varepsilon Z_t) \\
&= -\sigma_\varepsilon Z_{t+1} Z_t + \sigma_\varepsilon \chi^2_{1,t} \\
&= -\frac{\sigma_\varepsilon}{2}(\chi^2_{1,+} - \chi^2_{1,-}) + \sigma_\varepsilon \chi^2_{1,t}
\end{aligned}$$

$$E[K_{t+1,t}] = \sigma_\varepsilon, \quad \text{Var}[K_{t+1,t}] = 3\sigma_\varepsilon^2$$

, where we used the fact that the squared of a standard normal distribution Z is a chi-squared distribution χ^2_1 with 1 degree of freedom. The product of two independent standard normals can also be rewritten in terms of chi-square distribution.

$$\begin{aligned}
Z_a Z_b &= \frac{1}{4}(Z_a + Z_b)^2 - \frac{1}{4}(Z_a - Z_b)^2 \\
&= \frac{1}{2}(\chi^2_{1,+} - \chi^2_{1,-})
\end{aligned}$$

$$Z_a \pm Z_b \sim N(0,2), \quad (Z_a \pm Z_b)^2 \sim 2\chi^2_1$$

$$E[\chi^2_k] = k, \quad \text{Var}[\chi^2_k] = 2k$$

The expected profit is simply the entry volatility.

4.5.2 Unleveraged Strategy

The above leveraged/naive strategy is too risky to trade in practice; as we can see, the volatility of the profit is $\sqrt{3}$ times the expected profit. However, it is used to illustrate the idea of short volatility because the math is traceable and follows a proper distribution. In classical pairs trading, we enter the position only with $\text{sign}(\varepsilon_t)$ as the trade direction and filter out some of the signals using the z-score = $\varepsilon_t/\hat{\sigma}_\varepsilon$ with a proportional of the total capital, c. Here, we will assume $c = 1$ and enter all positions to simplify the math and highlight intuition; it, of course, generalizes to all filtering cases with various constraints.

$$
\begin{aligned}
q_t &= -\text{sign}(\varepsilon_t) \\
K_{t+1,t} &= q_t(\varepsilon_{t+1} - \varepsilon_t) \\
&= (\varepsilon_t - \varepsilon_{t+1})1_{[\varepsilon_t > 0]} + (\varepsilon_{t+1} - \varepsilon_t)1_{[\varepsilon_t < 0]} \\
&= (\varepsilon_t - \varepsilon_{t+1})[2 \cdot 1_{[\varepsilon_t > 0]} - 1]
\end{aligned}
$$

As we can see now, the marginal profit is a mixture of two truncated normal distributions with the following properties.

$$
\begin{aligned}
X &\sim N(\mu, \sigma^2)1(a < X < b) \\
E[X] &= \mu + \frac{\varphi(a) - \varphi(b)}{\Phi(b) - \Phi(a)}\sigma \\
\varphi(a) &= \frac{1}{\sqrt{2\pi}} \exp(-\frac{1}{2}a^2)
\end{aligned}
$$

$$
\Phi(a = \infty) = 1, \quad \Phi(a = -\infty) = 0, \quad \Phi(a = 0) = 0.5
$$

Finally, we can derive the expectation of the marginal profit, which shows that the unleveraged strategy is also a short volatility strategy.

$$
\begin{aligned}
E[K_{t+1,t}] &= E[(\varepsilon_t - \varepsilon_{t+1})1(\varepsilon_t > 0) + (\varepsilon_{t+1} - \varepsilon_t)1(\varepsilon_t < 0)] \\
&= E[\varepsilon_t 1(\varepsilon_t > 0) - \varepsilon_{t+1}1(\varepsilon_t > 0) + \varepsilon_{t+1}1(\varepsilon_t < 0) - \varepsilon_t 1(\varepsilon_t < 0)] \\
&= E[\varepsilon_t 1(\varepsilon_t > 0) - \varepsilon_t 1(\varepsilon_t < 0)] \\
&= E[N(X|0, \hat{\sigma}_\varepsilon^2)1(X > 0)] \quad \text{Gaussian symmetry} \\
&= \frac{2}{\sqrt{2\pi}}\hat{\sigma}_\varepsilon
\end{aligned}
$$

4.6 Conditional Profitability

Understanding the Pnl through the marginal is great for theory development; it tells us, on average, what is the risk-reward between any two pairs of ε. However, in practice, we are only able to invest in one realization of the underlying stochastic process. Hence, the conditional is much more important and informative for investors because one should expect, based on all kinds of different market conditions, to bet differently. In other words, if a time series had meander to the point ε_t, the knowledge of forward distribution $\varepsilon_{t+h}|\varepsilon_t$ should inform how one makes a decision right now. Such a dependence on trials is where opportunities are offered. If one can find an edge, she should invest. On the other hand, if the conditional distribution is symmetric with 0 means, the investor should not bet anything. It, of course, generalizes to multi-step predictions; simply change 1 to h, and then everything follows the usual predictive posterior distributions. Interestingly, we have a closed-form solution for the h-step predictive distribution for pairs trading previously derived. For the following, I will investigate the strategy along with the Kelly fraction.

$$
\begin{aligned}
q_t &= -\text{sign}(\varepsilon_t) \\
\varepsilon_{t+1}|\varepsilon_t, \theta &\sim N(\phi\varepsilon_t, \sigma_y^2(1-\rho^2)) \\
R_{t+1}|\varepsilon_t, \theta &= q_t \frac{\varepsilon_{t+1} - \varepsilon_t}{\varepsilon_t}|\varepsilon_t, \theta \\
&\sim N(q_t \frac{\phi\varepsilon_t - \varepsilon_t}{\varepsilon_t}, q_t^2 \frac{\sigma_y^2}{\varepsilon_t^2}(1-\rho^2)) \\
&\sim N(|1-\phi|, \frac{\sigma_y^2}{\varepsilon_t^2}(1-\rho^2)) \\
&\sim N(m = 1-\phi, s^2 = \frac{\sigma_y^2}{\varepsilon_t^2}(1-\rho^2)), \quad |\phi| < 1
\end{aligned}
$$

Here we transformed the predictive distribution to the nominal return R_t. Interestingly, we see the variance decreases the larger the $|\varepsilon_t|$. In addition, $s^2 = \frac{\sigma_y^2}{\varepsilon_t^2}(1-\rho^2)) = \frac{\sigma_\varepsilon^2}{\varepsilon_t^2} = 1/z_t^2$, the conditional variances can be thought as the inverse of instantaneous z-score squared. We know in pairs trading, we enter the unit position with $q_t = sign(\varepsilon_t)$ as the trade direction because this is where the expectations point to. However, it is only a placeholder; we also need to determine the strategy betting fraction. We use the previously derived continuous approximation of the Kelly fraction here.

$$
G(f)|\varepsilon_t = \sum_{t=1}^{T} \ln(1 + r + f_t(R_{t+1} - r))
$$

We can see the optimal Kelly fraction is clearly different for different times because of the conditionals. Let's just take one time step and consider the myopic one-step optimization, $G(f,t)|\varepsilon_t$. The expected log growth follows.

$$
\begin{aligned}
g_\infty(f) &= r + f(m-r) - \frac{1}{2}f^2 s^2 \\
&= r + f(1-\phi-r) - \frac{1}{2}f^2\frac{\sigma_y^2}{\varepsilon_t^2}(1-\rho^2) \\
\frac{\partial}{\partial f}g_\infty(f) &= (1-\phi-r) - f\frac{\sigma_y^2}{\varepsilon_t^2}(1-\rho^2) = 0 \\
f_t^* &= \frac{\varepsilon_t^2(1-\phi-r)}{\sigma_y^2(1-\rho^2)} = \frac{\varepsilon_t^2}{\sigma_\varepsilon^2}(m-r) \\
&= z_t^2(m-r) = \frac{m-r}{s^2}
\end{aligned}
$$

$$
\begin{aligned}
g_\infty(f^*) &= r + \frac{\varepsilon_t^2(1-\phi-r)}{\sigma_y^2(1-\rho^2)}(1-\phi-r) - \frac{1}{2}\frac{\varepsilon_t^4(1-\phi-r)^2}{\sigma_y^4(1-\rho^2)^2}\frac{\sigma_y^2}{\varepsilon_t^2}(1-\rho^2) \\
&= r + \frac{\varepsilon_t^2(1-\phi-r)^2}{2\sigma_y^2(1-\rho^2)} = r + \frac{(m-r)^2}{2s^2} \\
&= r + \frac{1}{2}[z(m-r)]^2 \\
&= r + \frac{1}{2}\tilde{S}^2 \\
V_\infty(f^*) &= \frac{\varepsilon_t^2(1-\phi-r)^2}{\sigma_y^2(1-\rho^2)} = \frac{(m-r)^2}{s^2} \\
&= [z(m-r)]^2 = \tilde{S}^2
\end{aligned}
$$

As usual, $\tilde{S}$ is the risk-free-rate-adjusted Sharpe ratio. The key takeaways: (1) The larger the instantaneous volatility of the synthetic portfolio $\sigma_\varepsilon^2 = \sigma_y^2(1-\rho^2)$, the smaller the betting size. (2) On the other hand, the smaller the pairs correlation ρ, the larger the betting size. (3) the smaller the ϕ, the faster the convergence rate, the larger the bet. (4) If the risk-free rate is larger than the mean return $1-\phi$, we should short the synthetic instead. (5) Lastly, the expected log-growth rate is proportional to the Sharpe $\tilde{S}$ of the synthetic portfolio that is not a function of the fraction. In other words, (5) can be effectively used to pre-evaluate the pair performance using only the ε_t trace.

4.7 Relationship with OU Process

For other necessary parameters, we can use the Ornstein–Uhlenbec (OU) process to model them. Mathematically, we can also model the ε_t as a mean reversion process:

$$d\varepsilon_t = k\,(m - \varepsilon_t)\,dt + \sigma dW_t$$

the analytic solution is:

$$\varepsilon_t = \left(1 - e^{-kt}\right) m + \varepsilon_0 e^{-kt} + \sigma \int_0^t e^{-k(t-s)} dW_s$$

the expectation is:

$$\begin{aligned} E\,[\varepsilon_{t+\tau} \mid \varepsilon_t] &= (1 - e^{-kt})m + \varepsilon_t e^{-kt} \\ &= m + (\varepsilon_t - m)e^{-k\tau} \end{aligned}$$

and the variance is:

$$\begin{aligned} \mathrm{Var}\,[\varepsilon_{t+\tau} \mid \varepsilon_t] &= V[\left(1 - e^{-k\tau}\right) m + \varepsilon_0 e^{-k\tau}] + V[\sigma \int_t^{t+\tau} e^{-k(t+\tau-s)} dW_s] \\ &= \sigma^2 \int_t^{t+\tau} V[e^{-k(t+\tau-s)}]ds, \quad dW \sim N(0,1) \\ &= \sigma^2 \int_t^{t+\tau} e^{-2k(t+\tau-s)}ds, \quad u = -2k(t+\tau-s) \\ &= \sigma^2 \int_{-2k\tau}^{0} e^u \frac{1}{2k} du \\ &= \sigma^2 \frac{1 - e^{-2k\tau}}{2k} \end{aligned}$$

To estimate the parameters, we use the $AR(1)$ model. First, we have to make a connection between the continuous process and the sample observation of it. Say, we have a linear function that the new observation X_t is a function of the previous observation X_{t-1}, which resembles a discrete Markov process.

$$X_t = a + bX_{t-1} + cZ_t, \quad Z_t \sim N(0,1)$$

$$= \mu - \phi\Delta t\mu + \phi\Delta t X_{t-1} + \sqrt{\sigma^2\Delta t}Z_t$$

$$= \mu + \phi(X_{t-1} - \mu)\Delta t + \sqrt{\sigma^2\Delta t}Z_t$$

$$a = \mu - \phi\Delta t\mu, \quad b = \phi\Delta t, \quad c = \sqrt{\sigma^2\Delta t}$$

We can see, with the sampling rate Δt and by substituting in some parameters $\{a,b,c\}$, X_t is exactly the Euler–Maruyama discretization of the OU process at time $\{k\Delta t\}_{k=0}^{\infty}$.

For estimation, we first compute the cumulative return over the entire lookback period

$$X_t = \sum_{i=\tau-T}^{t} \varepsilon_i, \quad t = \tau - h, \ldots, \tau$$

, then use the $AR(1)$ model to model the cumulative residuals $X_{\tau-h}, ..., X_\tau$

$$X_t = \hat{\mu} + \hat{\phi}X_{t-1} + \hat{\varepsilon}_{t+1}, \quad t = \tau - h, \ldots, \tau - 1$$

$$\hat{\varepsilon} \sim (0, \sigma_\varepsilon^2)$$

and we compare the regression result with the parameters

$$m = E\left[X_{eq}\right] = (1 - \hat{\phi}^h)\hat{\mu}, \quad \sigma_{eq}^2 = \text{Var}\left[X_{eq}\right] = \frac{1 - \hat{\phi}^{2h}}{1 - \hat{\phi}^2}\sigma_\varepsilon^2$$

where m is the mean of this reverting process, κ is the reversion speed, and σ_{eq} is the volatility of this process. σ_{eq} is used for determining the weight of each pair. These parameters are also useful in the pair selection process.

Lastly, we see the discrete convergence rate ϕ and the continuous convergence rate κ have a relationship in asymptotic,

$$\lim_{h\to\infty} \frac{1 - \phi^{2h}}{1 - \phi^2} = \lim_{h\to\infty} \frac{1 - e^{-2k\tau}}{2k}$$

$$1 - \phi^2 = 2k$$

$$k = \frac{1}{2}(1 - \phi^2)$$

, which is obvious from the fundamental theorem of calculus that the continuous rate is the midpoint of the discrete rate.

A. Conditional Multivariate Normal Theory

The Multivariate Normal and its conditional form [32] lay the mathematical foundation of the modern portfolio theory and dynamic linear model. The reason for the former is highly related to Markowitz's idea that a linear combination of portfolios reduces variances while preserving a target return; hence, the alluded asset diversification claim. On the other hand, the latter has more connection with Bayesian inference because normal has a beautiful yet powerful conjugate form that allows fast and efficient computation. The section focuses on the former, and the next section B concerns the Bayesian updates.

Suppose we have a vector $x \in R^p$ that can be partitioned down to two arbitrary chunks $x = [x_1, x_2]^\intercal$ where $x_1 \in R^{p_1}$, $x_2 \in R^{p_2}$, and $p = p_1 + p_2$. We assume x follows a Multivariate Normal Distribution, which we can write out the probabilistic form as follows.

$$x = \begin{bmatrix} x_1 \\ x_2 \end{bmatrix} \sim \text{MVN}(\begin{bmatrix} m_1 \\ m_2 \end{bmatrix}, V)$$

$$V = \begin{bmatrix} V_1 & R \\ R^\intercal & V_2 \end{bmatrix} \quad K = \begin{bmatrix} K_1 & H \\ H^\intercal & K_2 \end{bmatrix} = V^{-1}$$

$\{m_1, m_2\}$ are the mean vectors, V is the covariance matrix, and K is the corresponding precision matrix. Using the partition matrix and the relationship between the covariance and precision matrix, we can derive the following equations between the components. X and Y are temporary placeholders that we need to derive in order to decompose the matrix down to a triangular form, which is easier to do inversion.

$$
\begin{aligned}
V &= \begin{bmatrix} X & Y \\ 0 & I_{p_2} \end{bmatrix} \begin{bmatrix} I_{p_1} & 0 \\ R^\intercal & V_2 \end{bmatrix} \\
&= \begin{bmatrix} K_1^{-1} & RV_2^{-1} \\ 0 & I_{p_2} \end{bmatrix} \begin{bmatrix} I_{p_1} & 0 \\ R^\intercal & V_2 \end{bmatrix} \\
K_1^{-1} &= V_1 - RV_2^{-1}R^\intercal
\end{aligned}
$$

$$
\begin{aligned}
K = V^{-1} &= \begin{bmatrix} I_{p_1} & 0 \\ -V_2^{-1}R^\intercal & V_2^{-1} \end{bmatrix} \begin{bmatrix} K_1 & -K_1RV_2^{-1} \\ 0 & I_{p_2} \end{bmatrix} \\
&= \begin{bmatrix} K_1 & -K_1RV_2^{-1} \\ -V_2^{-1}R^\intercal K_1 & V_2^{-1}R^\intercal K_1 R V_2^{-1} + V_2^{-1} \end{bmatrix} \\
&= \begin{bmatrix} K_1 & H \\ H^\intercal & V_2^{-1} + H^\intercal K^{-1} H \end{bmatrix}
\end{aligned}
$$

Therefore, we can write the partition from the precision matrix in terms of partitions from the covariance matrix.

$$
\begin{aligned}
K_1^{-1} &= V_1 - RV_2^{-1}R^\intercal \\
K_2 &= V_2^{-1} + H^\intercal K_1^{-1} H \\
H &= -K_1RV_2^{-1}
\end{aligned}
$$

Then, suppose we know about x_2 and what the distribution for $x_1|x_2$ is the solution from the conditional distribution of the Multivariate Normal, which is also a Multivariate Normal distribution. The result can be easily derived using the Bayesian Theorem.

$$
\begin{aligned}
p(x_1|x_2) \propto\ & p(x_1, x_2) \\
& \propto \exp\left[-\frac{1}{2}(x-m)^\intercal K(x-m)\right] \\
& \propto \exp\left[-\frac{1}{2}[(x_1-m_1)^\intercal K_1(x_1-m_1)+(x_2-m_2)^\intercal K_2(x_2-m_2)-2(x_1-m_1)^\intercal H(x_2-m_2)]\right] \\
& \propto \exp\left[-\frac{1}{2}K_1[(x_1-m_1)^\intercal(x_1-m_1)-2(x_1-m_1)^\intercal K_1^{-1}H(x_2-m_2)]\right] \\
& \propto \exp\left[-\frac{1}{2}[x_1-(m_1-K_1^{-1}H(x_2-m_2))]^\intercal K_1[x_1-(m_1-K_1^{-1}H(x_2-m_2))]\right]
\end{aligned}
$$

$$
x_1|x_2 \sim \text{MVN}(m_1-K_1^{-1}H(x_2-m_2), K_1^{-1})
$$

The above result was established because the density follows the typical Multivariate Normal in proportionality. Further simplification, we can write $m_{1|2} = m_1 + A_{1|2}(x_2 - m_2)$ where $m_{1|2}$ is the conditional mean and $A_{1|2} = -K_1^{-1}H$ is usually referred to the regression matrix. Another interesting fact is the conditional covariance matrix is simply the inverse of the partitioned precision matrix, $V_{1|2} = K_1^{-1}$. This observation raises a more intriguing theory to the Gaussian Graphical Model that says the covariance matrix summarizes the information in terms of marginal while the precision matrix summarizes it in terms of conditional distributions. We will not investigate it here, but interested readers can refer to the citations for further information [9] [16] [6] [31] [14] [5]. To summarize the conditional normal results, we have the following.

$$x_1|x_2 \sim \text{MVN}(m_{1|2}, V_{1|2})$$

$$\begin{aligned}
m_{1|2} &= m_1 + A_{1|2}(x_2 - m_2) \\
A_{1|2} &= -K_1^{-1}H = -K_1^{-1}(-K_1 R V_2^{-1}) \\
&= R V_2^{-1}
\end{aligned}$$

$$\begin{aligned}
V_{1|2} &= K_1^{-1} \\
&= V_1 - R V_2^{-1} R^{\mathsf{T}}
\end{aligned}$$

B. Dynamic Linear Model

The system we deal with in finance and economics is a high-dimensional dynamic system that feeds and reacts consistently to new information. Any model that attempts to carry out inference or predictions on such a system based on static parameters is foolhardy. Because the moment the marketplace ingests fresh information, the trained static model would be rendered useless. Therefore, we need a model that has parameters that vary with time. Kalman Filter is particularly suited to such a problem and has been used in engineering science and control for many decades. It is a canonical example in the linear control theory literature. Here, I would like to introduce the Kalman Filter. A class of Bayesian filter that assumes conditional Gaussian structure and is under the broad umbrella of the Bayesian Dynamic Linear Model (DLM) [32] or, more generally, the State-Space model. Viewing the Kalman Filter through the lenses of Bayesian statistics makes the entire formulation much more intuitive and amenable than the traditional approach, which is how I will approach this. The beauty of this algorithm is that it is fully online and provides robust uncertainty estimation for a wide range of things. It is often found to be a component of a larger system of interest.

The general form of the DLM follows the below formulation.

$$\begin{aligned}
\text{Observation equation:} &\quad Y_t = F_t'\theta_t + \nu_t, \quad \nu_t \sim N(0, V_t) \\
\text{System equation:} &\quad \theta_t = G_t\theta_{t-1} + \omega_t, \quad \omega_t \sim N(0, W_t) \\
\text{Priors:} &\quad (\theta_0 | D_0) \sim N(m_0, M_0)
\end{aligned}$$

It consists of two parts: observations, Y_t, and the latent process θ_t. Any of them can be a vector that is indexed by time t. Then, we assume two independent error terms ν_t and ω_t, one for the observation equation and one more for the system/evolution equation. Often, ω_t is called innovation because it governs the assumption of how we believe the latent variable θ_t moves through time. F_t is the design matrix or the feature matrix, and G_t is the state transition matrix. Essentially, F_t maps the latent variable $theta_t$ to the observation Y_t, and G_t governs how the latent variable transitions through time. The two matrices are often pre-defined in the modeling step. There are ways to learn them, but it is out of the scope of this article because when we hone into a particular problem, the two matrices often become quite clear-cut and obvious. For the rest, we assume the two equations follow conditional Gaussian distributions with time-varying parameters. Therefore, we view them using the more familiar Bayesian notation below with the observation and transition densities with the parameter tuple $\{F_t, G_t, V_t, W_t\}$.

$$Y_t|\theta_t \sim N(F_t'\theta_t, V_t)$$

$$\theta_t|\theta_{t-1} \sim N(G_t\theta_{t-1}, W_t)$$

$$\theta_0 \sim N(m_0, M_0)$$

B.1 Constant First-Order Univariate Model

Now, let's consider the simplest model to highlight the most important concepts and features of the DLM with the parameter tuple $\{F_t = 1, G_t = 1, V_t = V, W_t = W\}$ under the assumption that we know both the observation variances and the innovation. No more vectors; let's just consider both Y_t and μ_t are one-dimensional univariate time series. Hence, it has the following form.

$$\begin{aligned} \text{Observation equation:} &\qquad Y_t = \mu_t + \nu_t, \quad \nu_t \sim N(0,V) \\ \text{System equation:} &\qquad \mu_t = \mu_{t-1} + \omega_t, \quad \omega_t \sim N(0,W) \\ \text{Priors:} &\qquad (\mu_0|D_0) \sim N(m_0, M_0) \end{aligned}$$

Here, we simply model the latent process μ_t as the mean of the time series that follows a purely stochastic process / random walk that is governed by the innovation ω_t. Then, the data we see is the mean plus some independent Gaussian noise. It is purely an empirical model and is often compared with other methods like the simple moving average or exponential moving average. The random walk assumption is simple and easy and is often used first in the exploratory phase. However, it basically assumes the underlying process is purely stochastic and has no value in prediction, modeling direction, or anything else. Often, this is not the case, but the random walk formulation does offer a simple starting point for many cases.

Since we assume V and W are known, it is natural to think of their ratio; it is called the signal-to-noise ratio (SNR) in engineering. W is the innovative/interesting part of the series, and V is the observation error that we would like to filter out. Later, we will generalize the model and consider the case when both V and W are unknown and infer them through data directly in an online setting. Now, we focus on how to update our beliefs when new information arrives so we can get the simplest model working as an example.

B.2 Update Equations: First-Order Univariate Model

Updating our belief is fairly straightforward because of the independent conditional Gaussian formulation in the model. We started with a prior belief $(\mu_0|D_0) \sim N(m_0, M_0)$, which is just a normal distribution with a mean and variances. Then, we observed many data $D_{t-1} = Y_1, Y_2, \ldots Y_{t-1}$. Our new belief stays in Gaussian but with updated parameters $(\mu_{t-1}|D_{t-1}) \sim N(m_{t-1}, M_{t-1})$. This is the starting point before you see new data Y_t comes in. Now, we will demonstrate how we incorporate the new information below. The procedure is called the Kalman predict-update step.

$$
\begin{aligned}
\text{Posterior for } \mu_{t-1}\text{:} \quad & \mu_{t-1}|D_{t-1} \sim N(m_{t-1}, M_{t-1}) \\
\text{Prior for } \mu_t\text{:} \quad & \mu_t|D_{t-1} \sim N(m_{t-1}, R_t = M_{t-1} + W_t) \\
\text{1-step forecast:} \quad & Y_t|D_{t-1} \sim N(f_t = m_{t-1}, Q_t = R_t + V_t) \\
\text{Posterior for } \mu_t\text{:} \quad & \mu_t|D_t \sim N(m_t, M_t) \\
\text{with} \quad & \quad m_t = m_{t-1} + A_t e_t, \quad M_t = A_t V_t \\
\text{with} \quad & \quad A_t = R_t/Q_t, \quad e_t = Y_t - f_t
\end{aligned}
$$

This completes one update cycle. The proof of the updating equations is trivial because the normal distribution is conjugate to normal; hence, the posterior has to be normal, and the result can be simply read off. We can interpret A_t just like the signal-to-noise ratio because it is a division between $R_t/Q_t = R_t/(R_t + V_t)$. R_t is the latent process's uncertainty, and $R_t + V_t$ is the expected variance of the observation. e_t is, of course, the one-step-ahead prediction error. Then, we update the mean m_t equal to the mean from last time plus some proportion of the prediction error, just like what we see in classical TD learning or gradient-based methods.

B.3 Discounting

Continue with the interpretation of A_t as the signal-to-noise ratio. The signal part is $R_t = M_{t-1} + W_t, M_{t-1}$ is the last step variance of the mean, and W_t is the added noises at this time step.

It is natural to think W_t leads to an increase of uncertainty of the underlying mean with time going by or loss of information. Hence, we can reparametrize $R_t = M_{t-1}/\delta$, with a constant δ that serves as an information discount rate. This allows us not to specify W_t directly, which is hard to come by in practice, but think of it in terms of a proportion of the observation noises V_t instead. In practice, I found using a value between 0.95 and 1 for the discount rate works quite well.

B.4 Unknown Variances

Now, we don't need to specify W_t anymore but consider the discount rate. We can also use the data to infer V_t, too. To simplify things, we consider the constant V in our context just to illustrate the idea. Under the Bayesian conjugate theory, we can impose a particular structure on W_t sequence and put a prior on V to ensure everything has analytical form as follows. It is essentially the normal-inverse-gamma conjugacy.

$$\begin{aligned}
\text{Observation equation:} \quad & Y_t = \mu_t + \nu_t, \quad \nu_t|\phi \sim N(0,V) \\
\text{System equation:} \quad & \mu_t = \mu_{t-1} + \omega_t, \quad \omega_t|\phi \sim N(0,VW_t^*) \\
\text{Priors:} \quad & (\mu_0|\phi,D_0) \sim N(m_0,VM_0^*) \\
& (\phi|D_0) \sim \text{Gamma}(n_0/2,(n_0 s_0)/2), \quad \phi = 1/V
\end{aligned}$$

B.5 Summary: First-Order Univariate Model

Therefore, coupling the concept of discounting and unknown observational variances prior, we arrive at the following updating expressions for the first-order univariate model, where the model definition was defined in the above sections.

$$
\begin{aligned}
\text{Posterior for } \mu_{t-1}: &\qquad \mu_{t-1}|\phi, D_{t-1} \sim N(m_{t-1}, VM^*_{t-1}) \\
\text{Prior for } \mu_t: &\qquad \mu_t|\phi, D_{t-1} \sim N(m_{t-1}, VR^*_t) \\
\text{with} &\qquad \quad R^*_t = M^*_{t-1}/\delta \\
\text{1-step forecast:} &\qquad Y_t|\phi, D_{t-1} \sim N(f_t = m_{t-1}, VQ^*_t) \\
\text{with} &\qquad \quad Q^*_t = R^*_t + 1 \\
\text{Posterior for } \phi_t: &\qquad \phi|D_t \sim \text{Gamma}(n_t/2, (n_t s_t)/2) \\
\text{with} &\qquad \quad s_t = (\delta n_{t-1} s_{t-1} + e^2_t/Q^*_t)/n_t \\
\text{with} &\qquad \quad n_t = \delta n_{t-1} + 1, \quad e_t = Y_t - f_t \\
\text{Posterior for } \mu_t: &\qquad \mu_t|\phi_t, D_t \sim N(m_t, VM^*_t) \\
\text{with} &\qquad \quad m_t = m_{t-1} + A_t e_t, \quad M^*_t = A_t \\
\text{with} &\qquad \quad A_t = R^*_t/Q^*_t
\end{aligned}
$$

This model has only one parameter to specify δ and 4 hyperparameters for the priors n_0, s_0, m_0, M^*_0, which the effect goes away quickly with more data. Using proper priors also provides robustness to cold start problems when, sometimes, a flight has not had any data yet. Lastly, because of the conjugacy of normal and gamma, we can marginalize ϕ out of the conditionals. The resulting mean and observations distributions follow Student T distributions, offer robustness on outliers, and serve as a good solution to cold starts.

$$
\begin{aligned}
\mu_t|D_t &\sim T_{n_t}(m_t, s_t M^*_t) \\
Y_t|D_t &\sim T_{n_t}(m_t, s_t Q^*_t)
\end{aligned}
$$

Bibliography

[1] Marco Avellaneda and Jeong-Hyun Lee. "Statistical arbitrage in the US equities market". In: *Quantitative Finance* 10.7 (2010), pp. 761–782. DOI: 10.1080/14697680903124632. eprint: https://doi.org/10.1080/14697680903124632. URL: https://doi.org/10.1080/14697680903124632.

[2] Fischer Black and Myron Scholes. "The Pricing of Options and Corporate Liabilities". In: *Journal of Political Economy* 81.3 (1973), pp. 637–654. ISSN: 00223808, 1537534X. URL: http://www.jstor.org/stable/1831029 (visited on 12/14/2023).

[3] L. BREIMAN. "OPTIMAL GAMBLING SYSTEMS FOR FAVORABLE GAMES". In: *The Kelly Capital Growth Investment Criterion*. Chap. 5, pp. 47–60. DOI: 10.1142/9789814293501_0005. eprint: https://www.worldscientific.com/doi/pdf/10.1142/9789814293501_0005. URL: https://www.worldscientific.com/doi/abs/10.1142/9789814293501_0005.

[4] João Caldeira and Guilherme Moura. "Selection of a Portfolio of Pairs Based on Cointegration: A Statistical Arbitrage Strategy". In: *Revista Brasileira de Finanças* 11 (Jan. 2013). DOI: 10.2139/ssrn.2196391.

[5] Carlos M. Carvalho, Hélène Massam, and Mike West. "Simulation of hyper-inverse Wishart distributions in graphical models". In: *Biometrika* 94.3 (Aug. 2007), pp. 647–659. ISSN: 0006-3444. DOI: 10.1093/biomet/asm056. eprint: https://academic.oup.com/biomet/article-pdf/94/3/647/733549/asm056.pdf. URL: https://doi.org/10.1093/biomet/asm056.

[6] Carlos M. Carvalho and Mike West. "Dynamic matrix-variate graphical models". In: *Bayesian Analysis* 2.1 (2007), pp. 69–97. DOI: 10.1214/07-BA204. URL: https://doi.org/10.1214/07-BA204.

[7] Mark Cummins and Andrea Bucca. "Quantitative Spread Trading on Crude Oil and Refined Products Markets". In: *Quantitative Finance* 12 (Aug. 2011). DOI: 10.2139/ssrn.1932471.

[8] Binh Do and Robert Faff. "Does Simple Pairs Trading Still Work?" In: *Financial Analysts Journal* 66.4 (2010), pp. 83–95. ISSN: 0015198X. URL: http://www.jstor.org/stable/25741293 (visited on 04/24/2023).

[9] Adrian Dobra et al. "Sparse graphical models for exploring gene expression data". In: *Journal of Multivariate Analysis* 90.1 (2004). Special Issue on Multivariate Methods in Genomic Data Analysis, pp. 196–212. ISSN: 0047-259X. DOI: https://doi.org/10.1016/j.jmva.2004.02.009. URL: https://www.sciencedirect.com/science/article/pii/S0047259X04000259.

[10] Robert Elliott, John van der Hoek, and William Malcolm. "Pairs Trading". In: *Quantitative Finance* 5 (June 2005), pp. 271–276. DOI: 10.1080/14697680500149370.

[11] Robert Engle. "GARCH 101: The Use of ARCH/GARCH Models in Applied Econometrics". In: *The Journal of Economic Perspectives* 15.4 (2001), pp. 157–168. ISSN: 08953309. URL: http://www.jstor.org/stable/2696523 (visited on 12/14/2023).

[12] Evan G Gatev, William N Goetzmann, and K. Geert Rouwenhorst. *Pairs Trading: Performance of a Relative Value Arbitrage Rule*. Working Paper 7032. National Bureau of Economic Research, Mar. 1999. DOI: 10.3386/w7032. URL: http://www.nber.org/papers/w7032.

[13] William Goetzmann, K. Rouwenhorst, and Evan Gatev. "Pairs Trading: Performance of a Relative Value Arbitrage Rule". In: *Review of Financial Studies* 19 (Feb. 2006), pp. 797–827. DOI: 10.1093/rfs/hhj020.

[14] Lutz Gruber and Mike West. "GPU-Accelerated Bayesian Learning and Forecasting in Simultaneous Graphical Dynamic Linear Models". In: *Bayesian Analysis* 11.1 (2016), pp. 125–149. DOI: 10.1214/15-BA946. URL: https://doi.org/10.1214/15-BA946.

[15] John C. Hull. *Options, futures, and other derivatives*. 6. ed., Pearson internat. ed. Upper Saddle River, NJ [u.a.]: Pearson Prentice Hall, 2006. XXII, 789. ISBN: 978-0-13-197705-1. URL: http://gso.gbv.de/DB=2.1/CMD?ACT=SRCHA&SRT=YOP&IKT=1016&TRM=ppn+563580607&sourceid=fbw_bibsonomy.

[16] Beatrix Jones and Mike West. "Covariance decomposition in undirected Gaussian graphical models". In: *Biometrika* 92.4 (Dec. 2005), pp. 779–786. DOI: 10.1093/biomet/92.4.779. eprint: https://academic.oup.com/biomet/article-pdf/92/4/779/1099549/924779.pdf. URL: https://doi.org/10.1093/biomet/92.4.779.

[17] J. L. Kelly Jr. "A New Interpretation of Information Rate". In: *Bell System Technical Journal* 35.4 (1956), pp. 917–926. DOI: https://doi.org/10.1002/j.1538-7305.1956.tb03809.x. eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/j.1538-7305.1956.tb03809.x. URL: https://onlinelibrary.wiley.com/doi/abs/10.1002/j.1538-7305.1956.tb03809.x.

[18] Nikhil Khanna. "The Kelly Criterion and the Stock Market". In: (2016). URL: https://sites.math.washington.edu/~morrow/336_16/2016papers/nikhil.pdf.

[19] Christopher Krauss. "STATISTICAL ARBITRAGE PAIRS TRADING STRATEGIES: REVIEW AND OUTLOOK". In: *Journal of Economic Surveys* 31.2 (2017), pp. 513–545. DOI: https://doi.org/10.1111/joes.12153. eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/joes.12153. URL: https://onlinelibrary.wiley.com/doi/abs/10.1111/joes.12153.

[20] B. Litterman and Q.R. Group. *Modern Investment Management: An Equilibrium Approach*. Wiley Finance. Wiley, 2004. ISBN: 9780471480655. URL: https://books.google.com/books?id=Fo0pjo7lXvcC.

[21] Harry Markowitz. "PORTFOLIO SELECTION*". In: *The Journal of Finance* 7.1 (1952), pp. 77–91. DOI: https://doi.org/10.1111/j.1540-6261.1952.tb01525.x. eprint: https:

//onlinelibrary.wiley.com/doi/pdf/10.1111/j.1540-6261.1952.tb01525.x. URL: https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1540-6261.1952.tb01525.x.

[22] Vasily Nekrasov. "Kelly Criterion for Multivariate Portfolios: A Model-Free Approach". In: *SSRN* (2014). DOI: http://dx.doi.org/10.2139/ssrn.2259133.

[23] Louis Rotando and Edward Thorp. "The Kelly Criterion and the Stock Market". In: *American Mathematical Monthly* 99 (Dec. 1992), pp. 922–931. DOI: 10.2307/2324484.

[24] William F. Sharpe. "Capital Asset Prices: A Theory of Market Equilibrium under Conditions of Risk". In: *The Journal of Finance* 19.3 (1964), pp. 425–442. ISSN: 00221082, 15406261. URL: http://www.jstor.org/stable/2977928 (visited on 02/12/2024).

[25] Jasper Snoek, Hugo Larochelle, and Ryan P. Adams. *Practical Bayesian Optimization of Machine Learning Algorithms*. 2012. arXiv: 1206.2944 [stat.ML].

[26] E. O. Thorp. "Optimal Gambling Systems for Favorable Games". In: *Revue de l'Institut International de Statistique / Review of the International Statistical Institute* 37.3 (1969), pp. 273–293. ISSN: 03731138. URL: http://www.jstor.org/stable/1402118 (visited on 02/05/2024).

[27] Edward Thorp. "A Perspective on Quantitative Finance: Models for Beating the Market". In: 2004. URL: https://api.semanticscholar.org/CorpusID:15199286.

[28] Edward Thorp. "Portfolio Choice and the Kelly Criterion". In: *Stoch. Models Finance* (Jan. 1971). DOI: 10.1016/b978-0-12-780850-5.50051-4.

[29] Edward Thorp. "The Kelly Criterion in Blackjack, Sports Betting, and the Stock Market". In: *Handbook of Asset and Liability Management* 1 (Dec. 2008). DOI: 10.1016/B978-044453248-0.50015-0.

[30] Ganapathy Vidyamurthy. "Pairs Trading: Quantitative Methods and Analysis / G. Vidyamurthy." In: (Jan. 2004).

[31] Hao Wang and Mike West. "Bayesian analysis of matrix normal graphical models". In: *Biometrika* 96.4 (Oct. 2009), pp. 821–834. ISSN: 0006-3444. DOI: 10.1093/biomet/asp049. eprint: https://academic.oup.com/biomet/article-pdf/96/4/821/587993/asp049.pdf. URL: https://doi.org/10.1093/biomet/asp049.

[32] Mike West and Jeff Harrison. "Bayesian Forecasting and Dynamic Models (2nd edn)". In: *Journal of the Operational Research Society* 49.2 (1998), pp. 179–179.