A Markowitz Approach to Managing a Dynamic Basket of Moving-Band Statistical Arbitrages

Kasper Johansson   Thomas Schmelzer   Stephen Boyd

June 8, 2024

Abstract

We consider the problem of managing a portfolio of moving-band statistical arbitrages (MBSAs), inspired by the Markowitz optimization framework. We show how to manage a dynamic basket of MBSAs, and illustrate the method on recent historical data, showing that it can perform very well in terms of risk-adjusted return, essentially uncorrelated with the market.

Contents

1 Introduction

We consider the problem of managing a portfolio of statistical arbitrages (stat-arbs), where each stat-arb is a mean-reverting portfolio of a subset of an n-asset universe. Trading strategies based on stat-arbs have been around since the 1980s and are popular due to their documented success in practice. The strategy is based on the idea that if the price of a portfolio of assets is mean-reverting, we can profit by buying the portfolio when it is cheap and selling it when it is expensive. The literature on stat-arbs is vast, and typically focuses on one of two aspects: (i) finding a linear combination of asset prices that is mean-reverting, or (ii) finding a trading policy that profits by exploiting the mean-reversion. To the best of our knowledge, there exists no satisfying solution to the problem of managing a portfolio of multiple stat-arbs, that is independent of the process of finding them. This paper aims to fill this gap and proposes a solution based on a Markowitz optimization approach. In particular, we focus on the case of moving-band stat-arbs (MBSAs), recently introduced in [JSB23], where the midpoint of the price band is allowed to move over time.

1.1 Related work

Stat-arbs. Stat-arb strategies have been popular ever since their introduction by Nunzio Tartaglia in the 1980s [Pol11, GGR06]. The first strategy was based on pairs trading, where the price difference between two assets is tracked and positions entered when this difference deviates from its mean. Its success has been demonstrated in numerous empirical studies in various markets like equities [AL10], commodities [Nak19, VM17], and currencies [FKD19]; see, *e.g.*, [GGR06, AL10, Per09, HJTW04, KDH17, CM13, Huc19, DGLR10]. The general setting of a stat-arb is a portfolio of (possibly more than two) assets that exhibits a mean-reverting behavior [FP16, §10.5]. The literature on stat-arbs tends to split into several categories: finding stat-arbs, modeling the (mean-reverting) portfolio price, and trading stat-arbs. For a detailed overview of the literature, we refer the reader to [Kra17].

MBSA. MBSAs were recently introduced as an extension of the tradition stat-arb, in which the midpoint of the price band is allowed to move over time [JSB23]. Although the method is new, it relies on principles related to Bollinger bands that have been known for decades [Bol92, Bol02]. The MBSA relies on solving a small convex optimization problem to find a portfolio of (possibly more than two) assets that vary within a moving band. The method has been shown to work well in practice, and is used in the empirical study in this paper.

Trading stat-arbs. With some exceptions, the literature on trading stat-arbs is mostly limited to pairs trading or trading individual stat-arbs, rather than managing a portfolio of several stat-arbs. Several methods are based on co-integration [Joh00, AD05, Gra83, EG87, Vid04]. For example, in [ZP18, ZZWP18, ZZP19] the authors consider a (non-convex) optimization problem for finding high variance, mean-reverting portfolios, in a co-integration space of several stat-arb spreads. Their strategy is based on finding a portfolio of spreads,

defined by a co-integration subspace, and implemented using sequential convex optimization. In [YP12a, PY18, YP18] the authors model the spread of a pair of assets as an autoregressive process, a discretization of the Ornstein-Uhlenbeck process, and show how to trade a portfolio of spreads under proportional transaction costs and gross exposure constraints using model predictive control. The authors of [YP12b] use co-integration techniques to find pairs of assets that are mean-reverting, and show how to construct optimal mean-variance portfolios of these pairs.

Our approach differs from the above methods in that we do not assume any particular model of the price process, and we do not use statistical analysis like co-integration tests to find stat-arbs. Instead, we assume several stat-arbs (defined below as a portfolio and a price signal) are given, and the problem is to find the optimal allocation to these stat-arbs. In other words, we decouple the problem of finding stat-arbs from the problem of portfolio construction.

1.2 Outline

The rest of this paper is structured as follows. In §2 we review MBSAs and set our notation. In §3 we show how to manage and evaluate a dynamic basket of MBSAs. We present an empirical study of the method on recent historical data in §4, and give conclusions in §5.

2 Moving-band stat-arbs

MBSAs were recently introduced in [JSB23]. We start with a universe of n assets, with $P_t \in \mathbf{R}_{++}^n$ denoting the price (suitably adjusted) in USD, in period $t = 1, 2, \ldots$.

2.1 Midpoint price and alpha

An MBSA is defined by a vector $s \in \mathbf{R}^n$ of asset holdings (in shares), with negative entries denoting short positions. The vector s is typically sparse, with only a modest number of nonzero entries, but that will not affect our method. The MBSA price in period t is $p_t = s^T P_t$. The MBSA *midpoint price* is given by the trailing M-period average price,

$$\mu_t = \frac{1}{M} \sum_{\tau=t-M+1}^{t} p_\tau, \tag{1}$$

where M is the rolling window memory. For a good MBSA, the difference of the price and midpoint price, $p_t - \mu_t$, will oscillate in a band; we obtain a profit by buying when the price difference is low and selling when it is high. We associate with the MBSA the 'alpha' value

$$\alpha_t = \mu_t - p_t. \tag{2}$$

If α_t is positive, we expect the price of the MBSA to increase, and if it is negative, we expect the price to decrease.

2.2 MBSA lifetime

In [JSB23] we show how to discover MBSAs (*i.e.*, the vector s) by approximately solving a constrained variance maximization problem, using the convex-concave procedure [SDGB16, LB16].

Each MBSA has a lifetime, starting at creation or discovery time $t = d$ and ending at time $t = e > d$, when we choose to decommission it. We say it is active over the periods $t = d, d+1, \ldots, e$. (See [JSB23, §2.3] for details.)

2.3 Multiple MBSAs

We consider a collection of multiple active MBSAs that changes over time as new ones are discovered and some existing ones are decommissioned. We denote K_t as the number of MBSAs active during period t, and index the MBSAs as $k = 1, \ldots, K_t$. The kth MBSA is defined by the holdings $s^{(k)} \in \mathbf{R}^n$, which gives rise to the price $p_t^{(k)} = (s^{(k)})^T P_t$, and alpha

$$\alpha_t^{(k)} = \mu_t^{(k)} - p_t^{(k)},$$

where $\mu_t^{(k)}$ is the moving midpoint of the kth MBSA. We collect the holdings of the K_t MBSAs in a matrix

$$S_t = [s^{(1)} \ \cdots \ s^{(K_t)}] \in \mathbf{R}^{n \times K_t},$$

and introduce the K_t-vectors

$$p_t = (p_t^{(1)}, \ldots, p_t^{(K_t)}), \qquad \mu_t = (\mu_t^{(1)}, \ldots, \mu_t^{(K_t)}), \qquad \alpha_t = (\alpha_t^{(1)}, \ldots, \alpha_t^{(K_t)}),$$

i.e., we extend the notation introduced in the previous paragraph to the setting of multiple MBSAs.

3 Managing a dynamic basket of MBSAs

3.1 Portfolio holdings and trades

Arb-level and asset-level portfolio holdings. The (MBSA) holdings during period t are denoted by $q_t \in \mathbf{R}^{K_t}$. (The units of q_t are 'shares' of the MBSAs.) The corresponding asset-level holdings are $h_t = S_t q_t \in \mathbf{R}^n$ (in shares), or, in USD, $P_t \circ (S_t q_t)$, where $\circ$ denotes the elementwise (Hadamard) product.

The cash account is denoted by $c_t \in \mathbf{R}$ (in USD). Since we will use the cash account as part of our collateral for short positions, we will assume $c_t > 0$ for all t, *i.e.*, we do not borrow cash. The total portfolio value at time t is then $p_t^T q_t + c_t$ in USD.

Trading cost. We denote the (asset level) trades vector at time t as $z_t = h_t - h_{t-1}$, *i.e.*, the change in asset-level holdings from time $t - 1$ to time t, in shares. The trading cost at time t is given by

$$(\kappa_t^{\text{trade}})^T |z_t|,$$

where $\kappa_t^{\text{trade}} \in \mathbf{R}_+^n$ is the vector of one half the bid-ask spread for each asset at time t, in USD per share, and the absolute value is elementwise.

Holding cost. The holding cost is given by

$$(\kappa_t^{\text{short}})^T(-h)_+,$$

where $\kappa_t^{\text{short}} \in \mathbf{R}^n$ is the vector of shorting rates for each asset over period t, in USD per share per trading period. Here $(u)_+ = \max\{u, 0\}$ is the nonnegative part of u, and in the expression above it is applied elementwise.

3.2 Markowitz objective and constraints

We manage a dynamic basket of MBSAs using a Markowitz optimization approach [Mar52, Mar59, BJK$^+$23], maximizing the expected return adjusted for transaction and holding costs, subject to constraints on the portfolio holdings which include a risk limit. We assume that we know the current asset prices P_t, the arb-to-asset transformation matrix S_t, and the previous portfolio holdings h_{t-1} (and q_{t-1}) and cash value c_{t-1}. Our goal is to find the new portfolio holdings h_t (and q_t and c_t). We denote the candidate values, which are optimization variables, as h, q, and c, respectively.

Objective. The objective is to maximize the alpha exposure, adjusted for transaction and shorting costs,

$$\alpha_t^T q - \gamma^{\text{trade}}(\kappa_t^{\text{trade}})^T(h - h_{t-1}) - \gamma^{\text{short}}(\kappa_t^{\text{short}})^T(-h)_+,$$

where γ^{trade} and γ^{short} are positive parameters trading off the alpha exposure, transaction cost, and holding cost. The true values of κ_t^{trade} and κ_t^{short} are not known at time t, but are estimated from historical data. This objective is a concave function of the variables h and q.

Cash-neutrality. We assume that the portfolio is cash-neutral [BJK$^+$23, GK00], *i.e.*,

$$p_t^T q = 0. \tag{3}$$

In other words, the total portfolio value is equal to the value of the cash account. (We can also add a limit on the market exposure (or beta), but we found empirically that this makes a small difference on top of the cash neutrality constraint.) This constraint is often used in long-short trading strategies [GK00, Chap. 15]. The cash neutral constraint is a linear equality constraint on the variable q.

Collateral constraint. We require that

$$P_t^T(h)_+ + (c)_+ \geq \eta(P_t^T(h)_- + (c)_-),$$

where $\eta \geq 1$ is a parameter. Here $(u)_- = (-u)_+ = \max\{-u, 0\}$ is the nonpositive part of u, applied elementwise to h above. This constraint ensures that the long position is at least η

times the short position. In the form above, it is not a convex constraint, but subtracting $(P_t^T(h_t)_- + (c_t)_-)$ from both sides yields the equivalent convex constraint

$$P_t^T h + c \geq (\eta - 1)(P_t^T(h)_- + (c)_-).$$

The cash-neutrality constraint implies $P_t^T h = 0$, so this simplifies to

$$c \geq (\eta - 1)P_t^T(h)_-,$$

i.e., the cash account is at least $(\eta - 1)$ times the short position. For example, we can set $\eta = 2.02$ to keep a collateral of 102% of the short position, as has typically been required by regulation [D'a02, GMR02].

MBSA size limit constraints. We add the constraints

$$|q|_k(p_t)_k \leq \xi_t^{(k)} c, \quad k = 1, \dots, K_t,$$

i.e., the absolute value of the position in the kth MBSA is at most $\xi_t^{(k)}$ times the total portfolio value (which is the same as the cash account), where $\xi_t^{(k)}$ is a parameter. These constraints serve two purposes. First, they act as a form of regularization in case there are more MBSAs than assets, in which case the covariance matrix in the risk constraint (defined below) will be ill-conditioned or singular. Second, they limit extreme positions in any single MBSA. During the active period $t = d, d+1, \dots, e$ of an MBSA, we set $\xi_t^{(k)} = \xi$ for $t = d, d+1, \dots, e-l$, and then linearly decrease $\xi_t^{(k)}$ to zero over the next l periods, where l is a parameter, say $l = 21$ periods, and ξ is a parameter common to all MBSAs. In other words we use the size limit constraints to gradually decommission MBSAs. These constraints are linear inequality constraints on the variables q and c, and so are convex.

Risk constraint. The traditional definition of risk of a portfolio is the variance of the portfolio return, expressed as a quadratic form of the asset weights with an estimated asset return covariance matrix. Taking the squareroot we obtain the standard deviation of portfolio return, i.e., its volatility, which we can convert to USD by multiplying by the portfolio value.

Here we propose a different risk measure, which we have found empirically to work better, and is more in line with the spirit of MBSAs. Our risk is based on fluctuation of the portfolio value around its expected midpoint, and is directly expressed in USD. Recall that the portfolio value is $p_t^T q_t$, which we expect to fluctuate around the midpoint value $\mu_t^T q_t$, both of these in USD. We take the risk to be an estimate of the short term mean square value of the difference of the portfolio value and the midpoint value, $(p_t - \mu_t)^T q_t$. We express this mean-square value (now using q, not q_t) as

$$q^T \Sigma_t q,$$

where Σ_t is an average of the recent values of $(p_\tau - \mu_\tau)(p_\tau - \mu_\tau)^T$. We limit our risk using the constraint

$$\|\Sigma_t^{1/2} q\|_2 \leq \sigma c,$$

7

where $\sigma > 0$ is a parameter setting the target risk (which is unitless) and c is the portfolio value. This risk limit is a convex constraint in the variables q and c, specifically a second-order cone (SOC) constraint [BV04, §4.4.2].

The covariance matrix Σ_t can be estimated in many ways; see, $e.g.$, [JOP+23]. We express it in terms of asset prices as

$$\Sigma_t = S_t^T \Sigma_t^P S_t,$$

where $\Sigma_t^P \in \mathbf{R}^{n \times n}$ is an estimate of the short-term covariance matrix of the asset prices. We estimate Σ_t^P as follows. First, define the centered price vectors

$$\tilde{P}_t = P_t - \bar{P}_t, \quad t = 1, \ldots, T,$$

where $\bar{P}_t$ is the M-period rolling window mean of the asset prices. (This is the same M used to define the MBSA midpoints in (1).) Then, Σ_t^P is an average of the recent values of $\tilde{P}_\tau \tilde{P}_\tau^T$. If a linear estimator, like a rolling window or exponentially weighted moving average (EWMA), is used to estimate this expectation, then this is equivalent to estimating Σ_t directly using the same method, $i.e.$, to center the MBSA prices and then compute an average of the outer product of the centered prices. We will use the iterated EWMA predictor [JOP+23, BB22], to estimate the covariance matrix of the centered prices $\tilde{P}_t$. (This is not equivalent to estimating Σ_t directly with an iterated EWMA, but in practice very similar.)

3.3   Markowitz formulation

We assemble the objective and constraints described above. We take h_t, q_t, and c_t, as optimal values of h, q, and c in the optimization problem

$$
\begin{aligned}
\text{maximize} \quad & \alpha_t^T q - \gamma^{\text{trade}}(\kappa_t^{\text{trade}})^T(h - h_{t-1}) - \gamma^{\text{short}}(\kappa_t^{\text{short}})^T(-h)_+ \\
\text{subject to} \quad & h = S_t q, \quad c \geq (\eta - 1)P_t^T h_-, \quad c = c_{t-1} + p_t^T q_{t-1}, \\
& p_t^T q = 0, \quad |q|_k (p_t)_k \leq \xi_t^{(k)} c, \quad k = 1, \ldots, K_t, \\
& \|\Sigma_t^{1/2} q\|_2 \leq \sigma c,
\end{aligned}
\tag{4}
$$

with variables $h \in \mathbf{R}^n$, $q \in \mathbf{R}^{K_t}$, and c. The problem data are:

- h_{t-1}, q_{t-1}, c_{t-1}, the previous portfolio holdings at the asset-level and arb-level, and cash;

- κ_t^{trade} and κ_t^{short}, (predictions of) the trading and holding costs;

- S_t, the arb-to-asset transformation matrix;

- η, the collateral parameter;

- P_t and p_t, the asset and MBSA prices;

- $\xi_t^{(k)}$, the MBSA size limits;

- Σ_t, a prediction of the short-term covariance matrix of the MBSA prices;

8

- σ, a target risk, expressed as a fraction of the portfolio value.

The problem (4) is a convex optimization problem, more specifically one that can be transformed to a second-order cone program (SOCP). The solution to this optimization problem gives the new MBSA portfolio q_t, the asset-level portfolio holdings h_t, and the cash account c_t at time t.

Extensions and variations. We can add additional constraints, such as a leverage constraint, asset position limits, maximum market exposure, *etc.* [BJK$^+$23]. We can also soften some constraints, if it is not critical that they are satisfied exactly and softening improves performance [BJK$^+$23]. For example, to soften the arb-to-asset constraint, we remove the constraint $h = S_t q$ and add a penalty term $\gamma^{\text{hold}} \| P_t \circ (h - S_t q) \|_1$ to the objective. (Note that softening the arb-to-asset constraint implicitly also softens the risk and cash-neutrality constraints if these are expressed in terms of the arb-level portfolio q.) Softening some constraints allows the optimizer to choose values that violate the original hard constraints when necessary, which can reduce unnecessary trading, and therefore transaction cost.

4 Numerical experiments

We illustrate the MBSA portfolio management strategy on recent historical data. Code to replicate the experiments is available at

$$\texttt{https://github.com/cvxgrp/cvxstatarb}.$$

4.1 Data and parameters

Data set. We gather daily price data from the CRSP US Stock Databases using the Wharton Research Data Services (WRDS) portal [WRD23]. The data set consists of adjusted asset prices of 15405 stocks from January 4, 2010, to December 30, 2023, for a total of 3282 trading days.

Monthly search for MBSAs. We search for MBSAs every 21 trading days, with the same setup and parameters as described in detail in [JSB23, §4.1].

Dynamic management of the MBSAs. Every day we solve the Markowitz optimization problem (4) to rebalance our portfolio. Each MBSA is kept in the portfolio for 500 trading days. After that ξ is reduced linearly to zero over the next 21 trading days.

Risk model. As described in §3.2, we decompose the covariance matrix as

$$\Sigma_t = S_t^T \Sigma_t^P S_t,$$

9

where Σ_t^P is the short-term covariance matrix of the asset prices. The covariance matrix Σ_t^P is estimated using the iterated EWMA (IEWMA) predictor (see [Eng02] and [JOP+23, Ch. 2.5]) on the centered prices $\tilde{P}_t = P_t - \bar{P}_t$, where $\bar{P}_t$ is the 21-day rolling window mean of the asset prices. For the IEWMA predictor, we use a 125-day half-life for volatility estimation, and a 250-day half-life for correlation estimation. To reduce trading induced by the risk-model, we smooth the covariances with a 250-day half-life EWMA [JOP+23, Chap. 9.2].

Parameters. For the MBSA problem we use the same parameters as in [JSB23, §4.1]. For the Markowitz optimization problem we use $\gamma^{\mathrm{trade}} = 1$, $\eta = 1$, $\xi = 1$, and $\sigma^{\mathrm{tar}} = 10\%$ annualized.

Trading and shorting costs. Our numerical experiments take into account transaction costs, *i.e.*, we buy assets at the ask price, which is the (midpoint) price plus one-half the bid-ask spread, and we sell assets at the bid price, which is the price minus one-half the bid-ask spread. We use 0.5% as a proxy for the annual shorting cost of stocks, which is well above what is typically observed in practice for liquid stocks [D'a02, GMR02, KL23, DD14]. We also note that we have tested the method with shorting costs upwards of 10% annually, and it remains profitable.

4.2 Simulation and metrics

This section explains how we simulate the portfolio, and the metrics we use to evaluate the performance.

Cash account. We initialize the portfolio with a cash account at time $t = 0$, $c_0 = 1$. The cash account evolves as

$$c_{t+1} = c_t - (q_{t+1} - q_t)^T p_{t+1} - \phi_t = c_t + q_t^T p_{t+1} - \phi_t, \quad t = 0, \dots, T,$$

where ϕ_t is the transaction and holding cost, consisting of the trading cost at time $t + 1$ and the holding cost over period t. (The last equality follows from the cash-neutrality constraint.)

Portfolio NAV. The net asset value (NAV) of the portfolio, including the cash account, at time t is

$$V_t = c_t + q_t^T p_t = c_t.$$

Due to the market neutrality constraint in (4) we have $V_t = c_t$, *i.e.*, the NAV is equal to the cash account.

Return. The return at time t is

$$r_t = \frac{V_t - V_{t-1}}{V_{t-1}}, \quad t = 1, \ldots, T.$$

We report several standard metrics based on the returns r_t. The average return is

$$\bar{r} = \frac{1}{T} \sum_{t=1}^{T} r_t,$$

which we multiply by 250 to annualize. The return volatility is

$$\left(\frac{1}{T} \sum_{t=1}^{T} (r_t - \bar{r})^2 \right)^{1/2},$$

which we multiply by $\sqrt{250}$ to annualize. The annualized Sharpe ratio is the ratio of the annualized average return to the annualized volatility. Finally, the maximum drawdown is

$$\max_{1 \le t_1 < t_2 \le T} \left(1 - \frac{V_{t_2}}{V_{t_1}} \right),$$

the maximum fractional drop in value form a previous high.

Turnover. The turnover at time t is

$$T = \frac{1}{2} \sum_{i=1}^{n} |((P_t \circ h_t)_i - (P_t \circ h_{t-1})_i)/V_t| = \frac{1}{2} \|P_t \circ (h_t - h_{t-1})\|_1 / V_t,$$

which we multiply by 250 to annualize. It measures the amount of trading in the portfolio [GK00, Chap. 16]. For example, a turnover of 0.01 means that the average of total amount bought and total amount sold is 1% of the total portfolio value.

Active return and risk. We define the active return as the return of the portfolio minus the return of the market, which we take to be the S&P 500. The active risk is the standard deviation of the active return.

Residual return and risk. Given the portfolio returns $r_1, r_2, \ldots, r_T$, and the corresponding market returns $r_1^m, r_2^m, \ldots, r_T^m$, we construct the linear model

$$r_t = \beta r_t^m + \theta_t, \quad t = 1, \ldots, T,$$

where βr_t^m is the return explained by the market (with $\beta \in \mathbf{R}^n$), and $\theta_t \in \mathbf{R}^n$ is the residual return at time t. The residual risk is the standard deviation of the residual return. The mean residual return is often referred to as the alpha of the portfolio [GK00] (not to be confused with the alpha of an MBSA). The information ratio is the ratio of the portfolio alpha to the residual risk.

Return	Volatility	Sharpe ratio	Turnover	Drawdown
19%	12%	1.61	136	15%

Table 1: Metrics for portfolio of MBSAs.

| Active return	Active risk	Residual return	Residual risk	Beta	Information ratio
8%	20%	18%	11%	11%	1.53

Table 2: Comparison of portfolio of MBSAs to the market.

4.3 Results

Portfolio performance. The portfolio performance is summarized in table 1. We attain an average annual return of 19% at an annual volatility of 12%, corresponding to a Sharpe ratio of 1.61, with a maximum drawdown of 15% over the roughly 10-year period. The average turnover is 136, which corresponds to a daily turnover of roughly 50%. In comparison to other stat-arb strategies in the literature, this seems like a reasonable level of turnover [GOPZ21].

Active and residual return and risk. Here we compare the MBSA strategy to the market, represented by the S&P 500 with an initial investment of $1 and diluted with cash to attain the same annualized risk as the MBSA strategy. Table 2 gives a numerical summary. The MBSA strategy attains an annualized Sharpe ratio of 1.61 (compared to 0.66 for the market) with a residual return (alpha) of 18% and a market beta of 11%. The information ratio is 1.53.

NAV evolution. Figure 1 shows the NAV of the MBSA strategy and the market, and the 250-day half-life EWMA correlation between the two. As seen, the MBSA strategy outperforms the market, with very low correlation, 15% over the whole period. This suggests mixing the MBSA strategy with the market. For example, mixing 90% of the MBSA strategy and 10% of the market yields an annualized Sharpe ratio of 1.66, slightly better than the MBSA strategy alone.
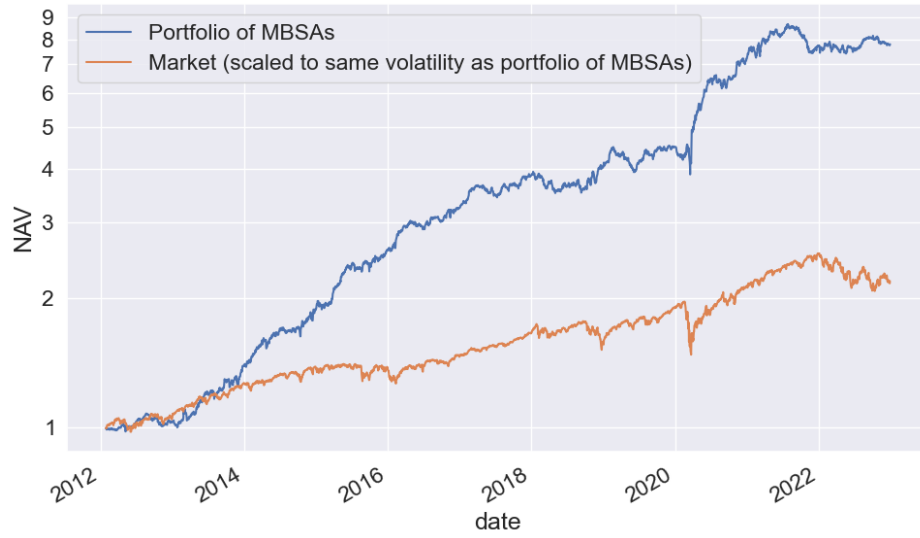
(a) NAV evolution.

(b) 250-day half-life EWMA correlation to market.

Figure 1: A comparison summary of the dynamically managed portfolio of MBSAs to the market.

Annual performance. Here we break down the metrics reported above over the 11-year period into the performance metrics for each year. Figure 2 shows the annual performance of the MBSA strategy and the market. The MBSA strategy has positive return in each of the 11 years, whereas the market has negative return in 2 of the 11 years. The MBSA strategy outperforms the market in 8 of the 11 years, and has more stable performance.

Finally, figure 3 shows the annual residual return and risk, and market beta. As seen, most of the MBSA success is unexplained by the market.
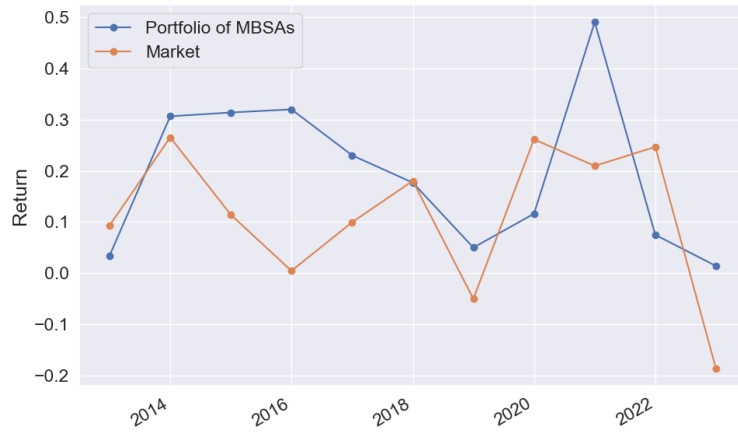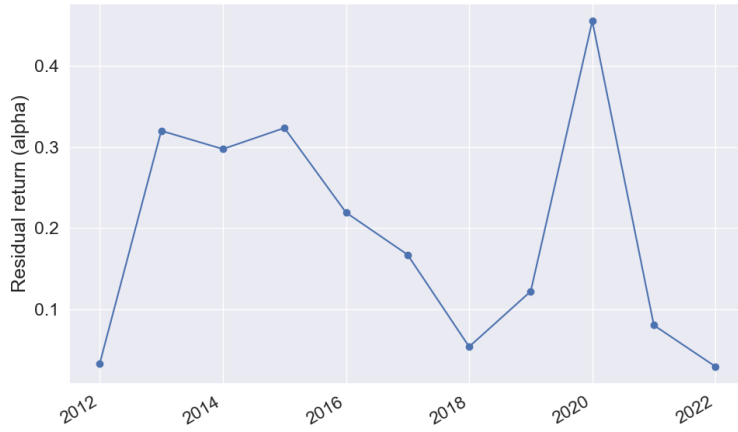
(a) Annual returns.
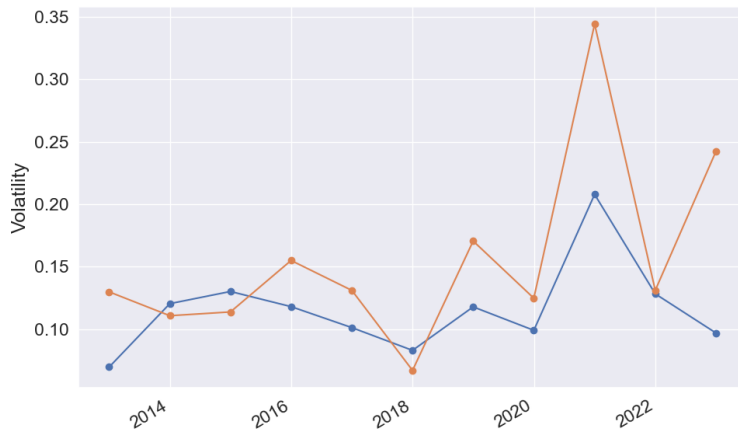
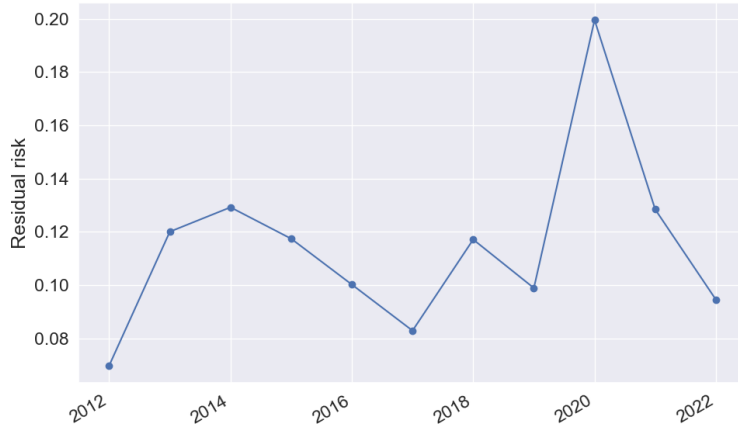(b) Annual volatilities.
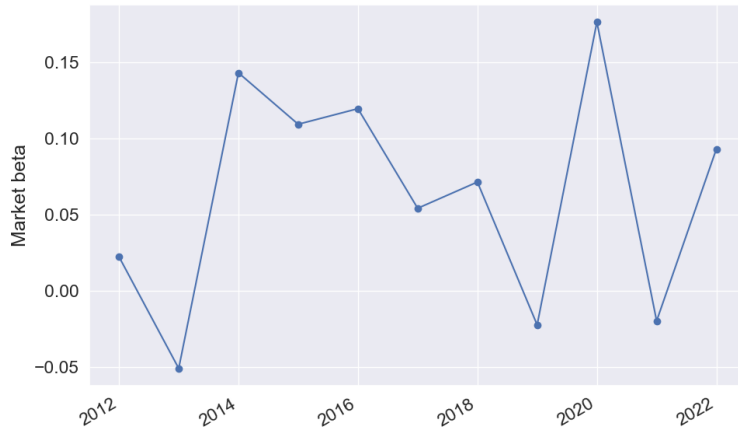

(c) Annual Sharpe ratios.

Figure 2: Annual portfolio performance.

(a) Annual residual returns (alphas).

(b) Annual residual risks.

(c) Annual market betas.

Figure 3: Annual residual return and risk, and market beta.

16

5 Conclusion

We have shown how to manage a dynamic basket of moving-band stat-arbs, based on a long-short Markowitz optimization strategy. We presented an empirical study of the method on recent historical data, showing that it can to outperform the market with low correlation.

References

[AD05]   C. Alexander and A. Dimitriu. Indexing and statistical arbitrage. *The Journal of Portfolio Management*, 31(2):50–63, 2005.

[AL10]   M. Avellaneda and J. Lee. Statistical arbitrage in the US equities market. *Quantitative Finance*, 10(7):761–782, 2010.

[BB22]   S. Barratt and S. Boyd. Covariance prediction via convex optimization. *Optimization and Engineering*, 2022.

[BJK⁺23] S. Boyd, K. Johansson, R. Kahn, P. Schiele, and T. Schmelzer. Markowitz portfolio construction at seventy, 2023. Working Paper, Stanford University.

[Bol92]   J. Bollinger. Using Bollinger bands. *Stocks & Commodities*, 10(2):47–51, 1992.

[Bol02]   J. Bollinger. *Bollinger on Bollinger bands*. McGraw-Hill New York, 2002.

[BV04]   S. Boyd and L. Vandenberghe. *Convex optimization*. Cambridge university press, 2004.

[CM13]   J. Caldeira and G. Moura. Selection of a portfolio of pairs based on co-integration: A statistical arbitrage strategy. *Available at SSRN 2196391*, 2013.

[D'a02]   G. D'avolio. The market for borrowing stock. *Journal of financial economics*, 66(2-3):271–306, 2002.

[DD14]   I. Drechsler and Q. Drechsler. The shorting premium and asset pricing anomalies. Technical report, National Bureau of Economic Research, 2014.

[DGLR10] C. Dunis, G. Giorgioni, J. Laws, and J. Rudy. Statistical arbitrage and high-frequency data with an application to Eurostoxx 50 equities. *Liverpool Business School*, 2010.

[EG87]   R. Engle and C. Granger. Co-integration and error correction: Representation, estimation, and testing. *Econometrica*, 55(2):251–276, 1987.

[Eng02]   R. Engle. Dynamic conditional correlation. *Journal of Business & Economic Statistics*, 20(3):339–350, 2002.

[FKD19]   T. Fischer, C. Krauss, and A. Deinert. Statistical arbitrage in cryptocurrency markets. *Journal of Risk and Financial Management*, 12(1):31, 2019.

[FP16]   Y. Feng and D. Palomar. A signal processing perspective on financial engineering. *Foundations and Trends® in Signal Processing*, 9(1–2):1–231, 2016.

[GGR06]   E. Gatev, W. Goetzmann, and G. Rouwenhorst. Pairs trading: Performance of a relative-value arbitrage rule. *The Review of Financial Studies*, 19(3):797–827, 2006.

[GK00]   R. Grinold and R. Kahn. *Active portfolio management*. McGraw Hill, 2000.

[GMR02]   C. Geczy, D. Musto, and A. Reed. Stocks are special too: An analysis of the equity lending market. *Journal of Financial Economics*, 66(2-3):241–269, 2002.

[GOPZ21]   J. Guijarro-Ordonez, M. Pelger, and G. Zanotti. Deep learning statistical arbitrage. *arXiv preprint arXiv:2106.04028*, 2021.

[Gra83]   C. Granger. *Co-integrated variables and error-correcting models*. PhD thesis, Discussion Paper 83-13. Department of Economics, University of California, 1983.

[HJTW04]   S. Hogan, R. Jarrow, M. Teo, and M. Warachka. Testing market efficiency using statistical arbitrage with applications to momentum and value strategies. *Journal of Financial economics*, 73(3):525–565, 2004.

[Huc19]   N. Huck. Large data sets and machine learning: Applications to statistical arbitrage. *European Journal of Operational Research*, 278(1):330–342, 2019.

[Joh00]   S. Johansen. Modelling of co-integration in the vector autoregressive model. *Economic Modelling*, 17(3):359–373, 2000.

[JOP+23]   K. Johansson, M. Ogut, M. Pelger, T. Schmelzer, and S. Boyd. A simple method for predicting covariance matrices of financial returns. *Foundations and Trends® in Econometrics*, 12(4):324–407, 2023.

[JSB23]   K. Johansson, T. Schmelzer, and S. Boyd. Finding moving-band statistical arbitrages via convex-concave optimization, 2023. Working Paper, Stanford University.

[KDH17]   C. Krauss, X. Do, and N. Huck. Deep neural networks, gradient-boosted trees, random forests: Statistical arbitrage on the S&P 500. *European Journal of Operational Research*, 259(2):689–702, 2017.

[KL23]   D. Kim and B. Lee. Shorting costs and profitability of long–short strategies. *Accounting & Finance*, 63(1):277–316, 2023.

[Kra17]   C. Krauss. Statistical arbitrage pairs trading strategies: Review and outlook. *Journal of Economic Surveys*, 31(2):513–545, 2017.

[LB16]   T. Lipp and S. Boyd. Variations and extension of the convex–concave procedure. *Optimization and Engineering*, 17:263–287, 2016.

[Mar52]   H. Markowitz. Portfolio selection. *The Journal of Finance*, 7(1):77–91, 1952.

[Mar59]   H. Markowitz. *Portfolio selection: Efficient diversification of investments.* John Wiley, 1959.

[Nak19]   T. Nakajima. Expectations for statistical arbitrage in energy futures markets. *Journal of Risk and Financial Management*, 12(1):14, 2019.

[Per09]   M. Perlin. Evaluation of pairs-trading strategy at the Brazilian financial market. *Journal of Derivatives & Hedge Funds*, 15:122–136, 2009.

[Pol11]   A. Pole. *Statistical arbitrage: algorithmic trading insights and techniques.* John Wiley & Sons, 2011.

[PY18]   J. Primbs and Y. Yamada. Pairs trading under transaction costs using model predictive control. *Quantitative Finance*, 18(6):885–895, 2018.

[SDGB16]   X. Shen, S. Diamond, Y. Gu, and S. Boyd. Disciplined convex-concave programming. In *2016 IEEE 55th Conference on Decision and Control (CDC)*, pages 1009–1014. IEEE, 2016.

[Vid04]   G. Vidyamurthy. *Pairs Trading: quantitative methods and analysis*, volume 217. John Wiley & Sons, 2004.

[VM17]   M. Vaitonis and S. Masteika. Statistical arbitrage trading strategy in commodity futures market with the use of nanoseconds historical data. In *Information and Software Technologies*, pages 303–313, Cham, 2017. Springer International Publishing.

[WRD23]   WRDS. Wharton research data services. `https://wrds-web.wharton.upenn.edu/wrds/`, 2023.

[YP12a]   Y. Yamada and J. Primbs. Model predictive control for optimal portfolios with co-integrated pairs of stocks. In *2012 IEEE 51st IEEE Conference on Decision and Control (CDC)*, pages 5705–5710. IEEE, 2012.

[YP12b]   Y. Yamada and J. Primbs. Optimal trading with cointegrated pairs of stocks. In *Recent Advances In Financial Engineering 2011*, pages 183–202. World Scientific, 2012.

[YP18]   Y. Yamada and J. Primbs. Model predictive control for optimal pairs trading portfolio with gross exposure and transaction cost constraints. *Asia-Pacific Financial Markets*, 25:1–21, 2018.

[ZP18]   Z. Zhao and D. Palomar. Mean-reverting portfolio with budget constraint. *IEEE Transactions on Signal Processing*, 66(9):2342–2357, 2018.

[ZZP19]   Z. Zhao, R. Zhou, and D. Palomar. Optimal mean-reverting portfolio with leverage constraint for statistical arbitrage in finance. *IEEE Transactions on Signal Processing*, 67(7):1681–1695, 2019.

[ZZWP18]   Z. Zhao, R. Zhou, Z. Wang, and D. Palomar. Optimal portfolio design for statistical arbitrage in finance. In *2018 IEEE Statistical Signal Processing Workshop (SSP)*, pages 801–805. IEEE, 2018.