Article

Semi-Metric Portfolio Optimization: A New Algorithm Reducing Simultaneous Asset Shocks

Nick James [1,*,†] , Max Menzies [2,†] and Jennifer Chan [3]

1 School of Mathematics and Statistics, University of Melbourne, Parkville, VIC 3010, Australia
2 Beijing Institute of Mathematical Sciences and Applications, Tsinghua University, Beijing 101408, China
3 School of Mathematics and Statistics, University of Sydney, Camperdown, NSW 2006, Australia
* Correspondence: nick.james@unimelb.edu.au
† These authors contributed equally to this work.

Abstract: This paper proposes a new method for financial portfolio optimization based on reducing simultaneous asset shocks across a collection of assets. This may be understood as an alternative approach to risk reduction in a portfolio based on a new mathematical quantity. First, we apply recently introduced semi-metrics between finite sets to determine the distance between time series' structural breaks. Then, we build on the classical portfolio optimization theory of Markowitz and use this distance between asset structural breaks for our penalty function, rather than portfolio variance. Our experiments are promising: on synthetic data, we show that our proposed method does indeed diversify among time series with highly similar structural breaks and enjoys advantages over existing metrics between sets. On real data, experiments illustrate that our proposed optimization method performs well relative to nine other commonly used options, producing the second-highest returns, the lowest volatility, and second-lowest drawdown. The main implication for this method in portfolio management is reducing simultaneous asset shocks and potentially sharp associated drawdowns during periods of highly similar structural breaks, such as a market crisis. Our method adds to a considerable literature of portfolio optimization techniques in econometrics and could complement these via portfolio averaging.

Keywords: portfolio optimization; time series analysis; change point detection; nonlinear dynamics; market crises

1. Introduction

One of the oldest and most important tasks in the field of econometrics is the analysis, forecasting, and optimization of financial risk. This may be conducted at the level of an individual stock, an entire sector, or a judiciously chosen portfolio. The most common measure of risk is portfolio variance, popularized by Markowitz (1952) in his seminal work. Markowitz' mathematical derivations assumed the Gaussianity of financial returns. Subsequently, returns of financial assets were shown to be non-Gaussian and fat-tailed in several works (Fama 1965; Mandelbrot 1963), prompting analysts to seek alternative measures of risk. Notably, tail risk measures such as value-at-risk (Braione and Scholtes 2016; Duffie and Pan 1997; Khraibani et al. 2018) or conditional value-at-risk/expected shortfall (Krause and Paolella 2014; Long et al. 2020; Tsay 2010; Ullah et al. 2022) have proven useful to guard against the greatest possible losses amid a financial crisis. In this paper, we propose an alternative approach to and measure of risk reduction, especially during a crisis, focusing on the diversification of assets away from simultaneous asset shocks, specifically in the form of coincident structural breaks.

Modern portfolio theory provides a framework for determining an allocation of weights in an investment portfolio by optimizing a specific objective function. The idea was first introduced by Markowitz (1952) and has progressed considerably since then. Markowitz' fundamental contribution was the concept of diversification among stock

portfolios, rather than analyzing risk and return on an individual security basis, one of the most seminal breakthroughs in econometrics. One of the most notable advancements was the work of Sharpe (1966), who proposed a measure of risk-adjusted returns in financial portfolios, the Sharpe ratio. This ratio is an indication of the potential reward in any candidate investment relative to its risk. The standard mathematical representation of the Sharpe ratio is the following optimization problem: given a collection of n assets, let R_i be the historical returns for the ith asset in a collection, Σ be the matrix of historical covariances between stocks, R_f the risk-free rate, and w_i the weights of the portfolio. One maximizes the Sharpe ratio, which we define by the following optimization problem:

$$\text{Maximize: } \frac{\sum_{i=1}^{n} w_i R_i - R_f}{\sqrt{\boldsymbol{w}^T \Sigma \boldsymbol{w}}}, \tag{1}$$

$$\text{subject to: } 0 \leq w_i \leq 1, i = 1, \dots, n, \tag{2}$$

$$\sum_{i=1}^{n} w_i = 1. \tag{3}$$

This objective function (1) selects an allocation of weights based on a trade-off between portfolio returns and variance. Returns are estimated from historical returns as $\mathbb{E}(R_p) = \sum_{i=1}^{n} w_i R_i$, while variance is estimated via $\sigma_p^2 = \boldsymbol{w}^T \Sigma \boldsymbol{w}$.

One may also impose certain conditions, depending on the context, which manifest as constraints, to accompany the objective function. The most common constraints, which we impose above and throughout the paper, are $0 \leq w_i \leq 1, i = 1, \dots, n$ (2) and $\sum_{i=1}^{n} w_i = 1$ (3). We also assume $R_f = 0$ throughout. These conditions require all portfolio assets to be invested and prohibit short selling, respectively. Weights are chosen by maximizing the objective function subject to such conditions. There is a wealth of other conditions that may be imposed, which is discussed further in the paper.

1.1. Overview of Portfolio Optimization

There has been significant research within the applied mathematics, computer science, and econometrics communities building upon Markowitz's mean-variance model (Markowitz 1952; Sharpe 1966). A variety of portfolio optimization frameworks have explored alternative objective functions utilizing risk measures other than standard volatility (Almahdi and Yang 2017; Bongini et al. 2002; Calvo et al. 2014; Soleimani et al. 2009; Vercher et al. 2007). Many authors in the field have taken existing theory and methodologies for the problem of portfolio selection and optimization (Bhansali 2007; Magdon-Ismail et al. 2003; Moody and Saffell 2001), including statistical mechanics (Li and Zhang 2021; Zhao and Xiao 2016), clustering (Iorio et al. 2018; León et al. 2017), fuzzy sets (Ammar and Khalifa 2003; Kocadağlı and Keskin 2015; Tanaka et al. 2000), graph theory (James et al. 2022), regularization (Fastrich et al. 2014; Li 2015; Pun and Wong 2019) regression trees (Cappelli et al. 2021), and multiobjective optimization (Lam et al. 2021; Mansour et al. 2019). For further details, a review of such techniques for portfolio optimization was conducted by Milhomem and Dantas (2020). Most importantly, for this paper, we must especially acknowledge the influence of statistical physics and econophysics. Many modern analyses of traditional financial markets (Alves et al. 2020; Eisler and Kertész 2006; James et al. 2022; Laloux et al. 1999), cryptocurrency markets (Drożdż et al. 2020, 2019, 2020, 2018; Pessa et al. 2023; Sigaki et al. 2019; Wątorek et al. 2021), and portfolio optimization (James et al. 2022; Prakash et al. 2021) have built upon methods developed in econometrics or inspired by physics (Basalto et al. 2007, 2008; Dose and Cincotti 2005; Drożdż et al. 2021; Fister et al. 2021; Gopikrishnan et al. 1998; James and Menzies 20212021a, c, 2022a, 2022b, 2022c, 2022d, 2022f, 2023a, 2023b; James et al. 2021, 2022, 2023; Liu et al. 1999; Mantegna et al. 2000; Podobnik et al. 2009; Valenti et al. 2018; Wang et al. 2006, 2020).

In particular, there has been a wealth of work that has modified the Sharpe ratio to penalize downside risk specifically. Sortino and van der Meer (1991) were pioneers in this, modifying the Sharpe ratio directly to only penalize downside variance. Since then, various

frameworks have been developed to directly target loss reduction, including value at risk models (Alexander and Baptista 2002; Campbell et al. 2001) and the mean–semivariance framework (Ballestero 2005; Boasson et al. 2011; Salah et al. 2018).

Finally, a substantial body of work has examined the difficulties provided by more complex investment constraints. Jin et al. (2016) provided a review of typical constraints in an asset allocation problem, as well as advances in algorithmic procedures (Liagkouras and Metaxiotis 2015, 2018; Lwin et al. 2014; Meghwani and Thakur 2017). In particular, cardinality constraints (Anagnostopoulos and Mamanis 2011) yield nonconvex sets (unions of lower-dimensional simplices) over which to perform optimization, providing a challenge to standard methods of convex optimization and producing NP-hard problems (Shaw et al. 2008).

In our paper, we compare our methodology primarily with those defined over long periods of training data. The primary reason we do this is to provide the algorithm adequate time to learn the systemic similarity in various asset's change point propagation. Given that change points indicate major shifts in underlying return dynamics, one can appreciate that these changes do not occur frequently. In fact, asset classes with especially low levels of beta (market risk) may only produce change points in the most extreme market conditions. Given that we wish to compute distances between change point propagation, which should include a sufficient number of change points, it is necessary that the training period is of significant length. That being said, our algorithmic approach may not need to be updated (and require the model being retrained) as frequently as other methods—as we are concerned with optimizing over very low frequency signals, which are unlikely to change quickly.

1.2. Overview of Change Point Detection Methods

Many domains in the physical and social sciences are interested in the identification of structural breaks in various data sets. Ranshous et al. (2015) and Akoglu et al. (2014) recently provided an overview of anomaly detection methods within the context of network analysis, which can be used to identify relations among entities in high-dimensional data. Koutra et al. (2016) determined change points (structural breaks) in dynamic networks via graph-based similarity measures, while James et al. (2021) analyzed change points in cryptocurrencies.

In the more econometric and statistical literature, focused on time series data, researchers have developed change point models driven by hypothesis tests, where p-values allow scientists to quantify the confidence in their algorithm (Bridges et al. 2015; Moreno and Neville 2013; Peel and Clauset 2015). Change point algorithms generally fall within statistical inference (namely Bayesian) or hypothesis testing frameworks. Bayesian change point algorithms (Adams and MacKay 2007; Barry and Hartigan 1993; James and Menzies 2022e; Xuan and Murphy 2007) identify change points in a probabilistic manner and allow for subjectivity through the use of prior distributions, but they suffer from hyperparameter sensitivity and do not provide statistical error bounds (p-values), often leading to a lack of reliability.

Within hypothesis testing, Ross (2015) outlined algorithmic developments in various change point models initially proposed by Hawkins (1977). Some of the more important developments in recent years include the work of Hawkins et al. (2003) and Ross (2014); Ross and Adams (2012); Ross et al. (2013). Ross (2015) recently created the CPM package, which allows for flexible implementation of various change point models on time series data. Given the package's ease of use, flexibility, and efficient implementation, we build our methodology on this suite of algorithms.

1.3. Overview of Semi-Metrics

The application of metric spaces has provided the groundwork for research advancement in various areas of machine learning. In addition to more traditional metrics, such as the Hausdorffand Wasserstein, semi-metrics, which may not satisfy the triangle inequality

property of a metric, have been used successfully in various machine learning applications. An overview of such (semi-)metrics and applications was recently provided by Conci and Kubrusly (2017). The three primary applications include image analysis (Baddeley 1992; Dubuisson and Jain 1994; Gardner et al. 2014), distance between fuzzy sets (Brass 2002; Fujita 2013; Gardner et al. 2014; Rosenfeld 1985), and computational methods (Atallah 1983; Atallah et al. 1991; Eiter and Mannila 1997; Shonkwiler 1989).

More recently, a review and computational analysis of various (semi-)metrics was undertaken by James et al. (2020) in measuring distance between time series' sets of structural breaks.

1.4. Motivation and Structure of This Paper

This paper aims to draw upon the aforementioned fields to yield a new approach to portfolio optimization with several benefits. While other existing methods aim to reduce downside risk directly (Boasson et al. 2011; Salah et al. 2018), we consider the significant shifts in asset behavior around structural breaks as a kind of "root cause" for simultaneous drawdown and paramount to avoid as the highest priority. Thus, we introduce the framework of using distances between structural breaks in our objective function; this necessitates distance measures between finite sets. In addition, we claim novelty in the precise methodology to measure discrepancy between finite sets. While previous research typically uses the Hausdorff or Wasserstein (Basalto et al. 2007, 2008) metric between sets, we use our own semi-metric, with favorable theoretical and empirical properties. Thus, our primary contributions in this paper are a novel framework of using structural breaks in portfolio optimization and a "proof of concept" via our specific implementation, as well as its validation. Validation is performed via simulated and real data, as well as a sequence of new propositions contrasting our chosen distance with previous options.

Our motivation in this paper is both theoretical and experimental. Theoretically, we investigate the use of distances between structural breaks in a broad attempt to highlight their potential utility in portfolio optimization applications, which we are unaware exists in the literature. We prove numerous properties of our particular choice of discrepancy between finite sets relative to existing alternatives. Additionally, experimentally, we offer a particular technique of portfolio optimization with promising results relative to existing options. The contribution of this paper goes beyond the specific methodology proposed; utilizing structural breaks may have numerous research directions. First, our specific methodology could be used within an ensemble framework, where our model could be one of several portfolio optimization procedures used to determine optimal portfolio weights (akin to model stacking and other ensemble-based methods). Second, our work could encourage researchers outside of econometrics and investment management specifically, but there is a general interest in allocating weights to various features, where each feature exhibits some sort of penalty in a time-varying fashion. This would be of particular interest in settings where the panel of data is especially stationary.

Thus, our paper is structured as follows. In Section 2, we outline our framework for optimization using structural breaks, including a detailed explanation of the general benefits of the framework. In Section 3, we explain the specific theoretical properties of our precise semi-metric compared with the existing Hausdorff or Wasserstein options to measure discrepancy between finite sets to explain the benefits of our precise choice. In Section 4.1, we use synthetic time series to show these benefits over other metrics more concretely in illustrated examples. Section 4.2 then performs a sample allocation of capital within a typical constrained optimization problem, featuring constraints frequently required by financial practitioners. In Section 5, we apply our method to real data across judiciously chosen training and testing periods, where we compare our methodology against nine others from the optimization literature. We conclude in Sections 6 and 7, summarizing the utility of our framework in the context of a typical asset allocation scenario.

2. Proposed Semi-Metric Change Point Optimization Framework

Our main contribution is to adapt the classical penalty function involving variance with one related to structural breaks. In many circumstances, variance is a suitable measure in a financial securities context. However, it is not without its limitations, and there are several reasons why a penalty function related to structural breaks may be a suitable alternative or complement to the covariance measure between two time series:

1. Covariance is computed as an expectation $\mathrm{Cov}(X, Y) = \mathbb{E}(X - \mathbb{E}X)(Y - \mathbb{E}Y)$, which is an average (integral) over an entire probability space. In a financial context, this computes an average over time; in modern financial markets, especially since the global financial crisis, most time periods are bull markets, with most assets performing quite well together. As such, assets that rise together in a bull market but actually exhibit distinct dynamics may be erroneously identified as similar.

2. Covariance fails to capture dissimilarity between time series during periods of market crisis and erratic behavior. Investors are often particularly concerned with the robustness of their portfolio during such times. Portfolios that are optimized using covariance as a risk measure fail to determine the impact of various asset combinations during times of market crisis. For instance, if two assets are simultaneously acting erratically, they may actually be negatively correlated during this time. If they are both included in a portfolio, this would increase rather than reduce erratic behavior. Structural breaks herald erratic behavior, so using distances between breaks in the objective function may better separate out erratic behavior in a portfolio.

3. Investors are also interested in peak-to-trough measures of asset performance, that is, the size of a drop in returns from a local maximum to a local minimum. Optimization algorithms using covariance measures fail to identify and minimize peak-to-trough behavior. However, distances between sets of structural breaks (in the mean, variance, and other stochastic quantities) are better equipped to identify how similar two time series are with respect to peak-to-trough measures. Thus, they may suitably allocate weights to minimize these precipitous drops.

4. While various methods of portfolio optimization target downside risk directly, we believe that structural breaks may be a kind of "root cause" of the greatest erratic behavior and simultaneous downside risk, and thus are of the greatest priority to diversify away from.

We formulate our new objective function to penalize structural breaks and their associated erratic behavior. We use the MJ_p family of semi-metrics of James et al. (2020). Given $p > 0$ and two nonempty finite sets $A, B \subset \mathbb{R}$ (or an arbitrary metric space), this is defined as

$$d_{MJ}^p(A, B) = \left(\frac{\sum_{b \in B} d(b, A)^p}{2|B|} + \frac{\sum_{a \in A} d(a, B)^p}{2|A|} \right)^{\frac{1}{p}}. \tag{4}$$

where $d(a, B)$ is the minimal distance from a point $a \in A$ to the finite set B. We note $d_{MJ}^p(A, B) = 0$ if and only if A, B. As discussed by James et al. (2020), varying $p > 0$ produces a family of semi-metrics, where larger values of p exhibit greater adherence to the triangle inequality property, but worse sensitivity to outliers. In our implementation, we select $p = 0.5$ due to its good performance with outlier sensitivity and the strong possibility of outliers in this context. Indeed, it is likely that some assets are impacted by market dynamics to which others are immune, which will yield outlier assets. We discuss distances between finite sets, their properties, and how we arrive at our family of semi-metrics in Appendix B.

We compute a distance matrix D_{ij} as follows: following a suitable change point algorithm (ours is described in Appendix A), let asset i have set of structural breaks $S_i, i = 1, \ldots, n$. Then, we form

$$D_{ij} = d_{MJ}^{0.5}(S_i, S_j). \tag{5}$$

In this paper, every set of structural breaks, simulated or real, is nonempty, so this computation is possible. Next, we transform our distance matrix into an affinity matrix, which mimics the properties of a covariance matrix:

$$A_{ij} = 1 - \frac{D_{ij}}{\max D}, \forall i, j. \tag{6}$$

Two assets have correlation equal to 1 if and only if they are perfectly correlated; analogously, $A_{ij} = 1$ if and only if $d(S_i, S_j) = 0$, meaning the two assets have identical structural breaks.

Remark 1. *We note that the denominator* $\max D$ *may be influenced by a single outlier time series. Thus, it is particularly important to choose our semi-metric (particularly the value of p) to handle outlier elements with care. We justify the benefits of choosing a smaller value of p, in this case* $p = \frac{1}{2}$, *in Proposition 3, Corollary 1, and Appendix B.*

In the context of Markowitz portfolio optimization, weights are chosen to maximize return while reducing total variance; this introduces more stocks with lower correlation, increases diversification, and reduces systematic risk in the portfolio. We modify this insight, allocating weights that maximize return while reducing affinity between sets of structural breaks, hence maximizing the spread between erratic behavior. To do so, we substitute our adjusted affinity matrix A for the original covariance matrix Σ and optimize a new risk-adjusted return measure with respect to portfolio weights. We term this the MJ ratio objective function (7), which we define in the following optimization problem:

$$\text{Maximize: } \frac{\sum_{i=1}^{n} w_i R_i - R_f}{\sqrt{w^T A w}}, \tag{7}$$

$$\text{subject to: } 0 \le w_i \le 1, i = 1, \dots, n, \tag{8}$$

$$\sum_{i=1}^{n} w_i = 1. \tag{9}$$

Essentially, this method retains the estimation of returns exactly as in the Sharpe ratio, $\mathbb{E}(R_p) = \sum_{i=1}^{n} w_i R_i$ and substitutes variance $\sigma_p^2 = w^T \Sigma w$ with a new denominator $\Omega_p^2 = w^T A w$, whose purpose is to "spread out" various assets' structural breaks.

Throughout the paper, we always retain at least the same constraints as Section 1, $0 \le w_i \le 1, i = 1, \dots, n$ (8) and $\sum_{i=1}^{n} w_i = 1$ (9). In subsequent sections, we also impose additional real-world constraints, such as upper and lower bounds on weights, and discuss how our method would work with other constraints frequently used in investment policy statements. Our method is flexible enough to vary such constraints, with no increase in complexity, provided such constraints result in a convex set of optimization. For more difficult constraints such as cardinality and preassignment constraints, our method could be combined with advances in the literature for efficient optimization over such spaces (Liagkouras and Metaxiotis 2018).

Remark 2. *One could also use this approach (substituting variance for affinity between structural breaks) to modify alternative existing portfolio selection methods, such as minimum variance optimization. Traditionally, this is defined as the following optimization problem:*

$$\text{Minimize: } \boldsymbol{w}^T \Sigma \boldsymbol{w}, \tag{10}$$

$$\text{subject to } \sum_{i=1}^{n} w_i R_i = P, \tag{11}$$

$$0 \leq w_i \leq 1, i = 1, \ldots, n, \tag{12}$$

$$\sum_{i=1}^{n} w_i = 1. \tag{13}$$

Typically, the desired level of returns P is selected according to the risk appetite of the investor. One can also incorporate a risk-free asset by including it as one of the permissible asset classes, with a term $R_0 = R_f$.

One could then write an equivalent of minimum variance optimization in our new context of structural breaks by formulating an optimization problem as follows:

$$\text{Minimize: } \boldsymbol{w}^T A \boldsymbol{w}, \tag{14}$$

$$\text{subject to } \sum_{i=1}^{n} w_i R_i = P, \tag{15}$$

$$0 \leq w_i \leq 1, i = 1, \ldots, n, \tag{16}$$

$$\sum_{i=1}^{n} w_i = 1. \tag{17}$$

Alternatively, one could formulate an analogy to the global minimum variance problem. This would be formulated simply by removing condition (15) as follows:

$$\text{Minimize: } \boldsymbol{w}^T A \boldsymbol{w}, \tag{18}$$

$$\text{subject to } 0 \leq w_i \leq 1, i = 1, \ldots, n, \tag{19}$$

$$\sum_{i=1}^{n} w_i = 1. \tag{20}$$

3. Theoretical Properties

In this section, we examine the mathematical properties of our proposed objective function and procedure and explain our choice of distance function between sets, including an analysis of alternatives.

Proposition 1. *The MJ ratio, as presented in (7), can be maximized on the chosen domain of weights, and the maximum can be determined analytically.*

Proof. First, we note that the matrix A is not necessarily positive semidefinite, so standard arguments regarding the optimization of the Sharpe ratio do not apply *mutatis muntandis* to the MJ ratio. Instead, we require a continuity and compactness argument. Due to the conditions on the weights, the ratio is optimized over a space $S = \{w_i : 0 \leq w_i \leq 1, \sum_{i=1}^{n} w_i = 1\}$. This is a compact space, specifically a $(n-1)-$simplex. By the definition of (6), all entries of A are non-negative, with diagonal elements equal to 1. Thus, $w^T A w$ is a continuous function on S that attains only positive values, and so the denominator of (7) is positive on the whole space S. This implies the MJ ratio is a well-defined continuous function on S. Since S is compact, it must achieve a global maximum on S.

Finally, since S is a $(n-1)-$simplex, one can examine and test the critical points within the simplex and use Lagrange multipliers on the boundary to find all possible

maxima and test them. In our implementation, we determine the optimal weights with a simple grid search. $\square$

Proposition 2 (Method complexity)**.** *Suppose we have n assets indexed $i = 1, ..., n$ over a time period of length T. Assume the weights have only the minimal constraints $0 \leq w_i \leq 1, \sum_{i=1}^{n} w_i = 1$. Then, the computational cost of the weight selection methodology described in Section 2 is $O(n^2 T + nT^2)$.*

Proof. As explained in Appendix A, the selection of change points for a single asset has a running time of $O(T^2)$ due to the two-phase procedure. Thus, the selection of change points $S_1, ..., S_n$ for all assets has cost $O(nT^2)$. For each i, j, the computation of $d_{MJ}^p(S_i, S_j)$ involves at most $O(|S_i| + |S_j|)$ comparisons between elements. Let $m = \max_i |S_i|$. This means the computation of each pairwise distance $d_{MJ}^p(S_i, S_j)$ is of complexity $O(m)$. Thus, the construction of the full distance matrix D and associated affinity matrix A is of complexity $O(mn^2)$. To select the weights, only the historical returns R_i and the matrix A are needed. We use sequential quadratic programming, which has a complexity cost of $O(n^2)$ when performed over a convex set with a fixed tolerance bound.

   Overall, the total cost of our procedure is $O(nT^2 + mn^2 + n^2)$; the three steps are each implemented in C++. We can use the simple bound $m \leq T$ to gain the final bound $O(n^2 T + nT^2)$. $\square$

Remark 3. *In our implementation, we consistently train our algorithm over long periods to appropriately learn relationships in structural breaks. Thus, it is usually the case that $n << T$. For example, a typical portfolio manager would have at most $n = 1000$ assets to choose from, while we train over $T = 2051$ days. Thus, the computational cost simplifies to $O(nT^2)$, with the T^2 operations implemented efficiently in C++. Thus, we can deduce that our method scales well with large numbers of stocks to choose from, with just a linear increase in complexity with the number of assets.*

   Furthermore, our complexity is unchanged with real-world upper and lower bounds on the weights $c_i \leq w_i \leq C_i$. These constraints, common in investment policy statements, still produce a convex and compact set over which we select the weights w_i, so the complexity is unchanged. As discussed in Section 5, these are the only additional constraints we impose in our experimentation, a common feature of real-world policy statements (Coffey 2016). When additional constraints are imposed, our optimization domain may be nonconvex. However, it will always be compact, so Proposition 1 holds. Efficient optimization of our objective function would require a combination with recent work in optimizing over domains with numerous nonconvex constraints imposed (Liagkouras and Metaxiotis 2015, 2018; Lwin et al. 2014; Meghwani and Thakur 2017; Shaw et al. 2008).

   In the following two propositions, we justify our selection of distance measure between finite sets, specifically two advantages it has over the popular Hausdorff and Wasserstein metrics between sets.

Definition 1 (Hausdorff metric)**.** *Let S, T be closed bounded subsets of $\mathbb{R}$ (or an arbitrary metric space). Their Hausdorff distance is defined by*

$$d_H(S, T) = \max \left(\sup_{s \in S} d(s, T), \sup_{t \in T} d(t, S) \right), \tag{21}$$

$$= \sup\{d(s, T), s \in S; d(t, S), t \in T\}, \tag{22}$$

where $d(s, T) = \inf_{t \in T} d(s, t)$ is the infimum distance from s to T.

   One could conceivably use this metric to measure distance between structural breaks, rather than our semi-metric. However, the Hausdorff metric suffers from substantial

sensitivity to outliers, as observed by Baddeley (1992). We formalize this in the following proposition:

Proposition 3. *Let* $T = \{t_1, \ldots, t_n\}$ *and* S *be fixed. Fixing all but one element, if* $t_n \to \infty$ *acts as an outlier, then the asymptotic behavior of the Hausdorff and* MJ_p *distances are as follows:*

$$d_H(S, T) \sim |t_n|, \qquad d_{MJ}^p(S, T) \sim \frac{|t_n|}{(2|T|)^{\frac{1}{p}}}, \tag{23}$$

$$i.e. \quad \lim_{|t_n| \to \infty} \frac{d_H(S, T)}{|t_n|} = 1, \lim_{|t_n| \to \infty} \frac{d_{MJ}^p(S, T)}{|t_n|} = \frac{1}{(2|T|)^{\frac{1}{p}}}. \tag{24}$$

Proof. For both $d_H(S, T)$ and $d_{MJ}^p(S, T)$, the only term that increases as $|t_n| \to \infty$ is $d(t_n, S)$, which increases asymptotically with $|t_n|$. The result follows immediately for d_H and follows for d_{MJ}^p by inspecting the coefficient $\frac{1}{2|T|}$ that accompanies the term $d(t_n, S)$. □

Corollary 1. *Let* $0 < p < q$ *and adopt the assumptions of Proposition 3. Then,* $d_H(S, T)$ *exhibits worse asymptotic outlier sensitivity than* d_{MJ}^q, *which itself is worse than* d_{MJ}^p.

Proof. If $p < q$, then $1 < (2|T|)^p < (2|T|)^q$. It follows that

$$\frac{1}{(2|T|)^{\frac{1}{p}}} < \frac{1}{(2|T|)^{\frac{1}{q}}} < 1. \tag{25}$$

Thus, the asymptotic coefficient of $|t_n|$ is the least for d_{MJ}^p, then d_{MJ}^q, and then d_H, as shown in (23) and (24). That is, a single element has the least influence on the increasing values of d_{MJ}^p than on d_{MJ}^q, and less so than d_H. □

As a consequence of this outlier sensitivity property, the Hausdorff metric d_H may grant an excessively high distance, and hence a low affinity, based on a single outlier structural break. In particular, two time series that have quite similar structural breaks (and hence erratic behavior profiles) may be granted low affinity and both be included in a portfolio based on just one structural break. On the other hand, the MJ_p semi-metric handles outliers well, and increasingly well with small values p, such as our choice in implementation $p = 0.5$. We illustrate an example of this in Section 4.1 and explain this further in Appendix B.

The other metric occasionally used to measure distance between finite sets is the Wasserstein metric. To be precise, it is most frequently employed between probability measures on a metric space, such as $\mathbb{R}$, as follows: let μ, ν be probability measures on $\mathbb{R}$, and $q \geq 1$, then

$$W_q(\mu, \nu) = \inf_\gamma \left(\int_{\mathbb{R} \times \mathbb{R}} |x - y|^q d\gamma \right)^{\frac{1}{q}}. \tag{26}$$

This infimum is taken over all joint probability measures γ on $\mathbb{R} \times \mathbb{R}$ with marginal probability measures μ and ν. The Formula (26) is difficult to compute in general, but in the case where μ, ν have cumulative distribution functions F, G on $\mathbb{R}$, there is a simple representation (del Barrio et al. 1999):

$$W_q(\mu, \nu) = \left(\int_0^1 |F^{-1} - G^{-1}|^q dx \right)^{\frac{1}{q}}, \tag{27}$$

where F^{-1} is the inverse cumulative distribution function, or more precisely, quantile function, associated to F (Gilchrist 2000). One can then use this to define a metric between

finite sets S, T. One associates to each set a probability measure defined as a weighted sum of Dirac delta measures

$$\mu_S = \frac{1}{|S|} \sum_{s \in S} \delta_s. \tag{28}$$

Then, the Wasserstein metric between sets S, T can be defined as $d_W^q(S, T) := W_q(\mu_S, \mu_T)$ and computed with (27). One could conceivably use this metric to measure distance between structural breaks instead of the Hausdorff metric. However, the Wasserstein metric has a property that makes it unsuitable in our context. Using the definition (28) and the Equation (27), the Wasserstein metric has a geometric property with respect to translation, $d_W^q(S, S + a) = |a|$. However, this is unsuitable for measuring the distance between sets with high intersection. We formalize these remarks in the following proposition:

Proposition 4. *If $|S \cap T| = r$, the following inequality holds:*

$$d_{MJ}^p(S, T) \le \left[1 - \frac{r}{2} \left(\frac{1}{|S|} + \frac{1}{|T|} \right) \right]^{\frac{1}{p}} d_H(S, T). \tag{29}$$

No such inequality holds for Wasserstein metric. Given a set S and its translation $S + a$ for some $a \in \mathbb{R}$, the Wasserstein metric has the property that $d_W^q(S, S + a) = |a|$. As a consequence, even with $|S \cap T| = |S| - 1 = |T| - 1$, it is possible for $d_W^q(S, T)$ to coincide with $d_H(S, T)$.

Proof. Examining the definition (4), any $d(s, T)$ or $d(t, S)$ term with $s \in S \cap T$ or $t \in S \cap T$, respectively, vanishes. Any other $d(s, T), d(t, S)$ term is at most $d_H(S, T)$. So,

$$d_{MJ}^p(S, T) \le \left[\frac{1}{2|S|} (|S| - r) + \frac{1}{2|T|} (|T| - r) \right] d_H(S, T), \tag{30}$$

which gives the inequality after simplifying. Turning to the Wasserstein metric, let $S = \{s_1, \dots, s_n\} \subset \mathbb{R}$ be a set with $s_1 < s_2 < \dots < s_n$ and $a \in \mathbb{R}$ a translate. Then, $S + a = \{s_1 + a, \dots, s_n + a\}$. By (27) and (28), $d_W^q(S, S + a)$ can be computed as

$$\left(\int_0^1 |F^{-1} - G^{-1}|^q dx \right)^{\frac{1}{q}}, \tag{31}$$

where F^{-1}, G^{-1} are the quantile functions associated to μ_S and μ_{S+a}. By integrating μ_S and μ_{S+a}, we can see that F, G are piecewise constant increasing step functions:

$$F = \sum_{j=1}^{n-1} \frac{j}{n} \mathbb{1}_{[s_j, s_{j+1})} + \mathbb{1}_{[s_n, \infty)}, \quad G = \sum_{j=1}^{n-1} \frac{j}{n} \mathbb{1}_{[s_j + a, s_{j+1} + a)} + \mathbb{1}_{[s_n + a, \infty)}. \tag{32}$$

It follows that their respective quantile functions are determined almost everywhere as

$$F^{-1} = \sum_{j=1}^{n} s_j \mathbb{1}_{\left(\frac{i-1}{n}, \frac{i}{n} \right)}, \quad G^{-1} = \sum_{j=1}^{n} (s_j + a) \mathbb{1}_{\left(\frac{i-1}{n}, \frac{i}{n} \right)}. \tag{33}$$

It follows quickly that $G^{-1} - F^{-1}$ is simply a constant function on $(0, 1)$ with value a, so the expression (31) simplifies to $|a|$. This concludes the second statement of the proposition.

Finally, let $S = \{0, 1, \dots, n-1\}$ and $T = \{1, 2, \dots, n\}$. As $T = S + 1$ is a translate of S, we have shown that $d_W^q(S, T) = 1 = d_H(S, T)$. However, $|T| = |S| = n$, while $|S \cap T| = n - 1$, showing that no such inequality as (29) holds for the Wasserstein metric. $\square$

As a consequence of this translation property, the Wasserstein metric d_W^q may excessively grant an excessively high distance, and hence a low affinity, to two sets with a very high intersection. For example, if two sets of structural breaks are $A = \{100, 200, \dots, 900\}$

and $B = \{200, 300, \ldots, 1000\}$, then d^p_{MJ} will reflect the high intersection and similarity between the sets A and B, while d^q_W will not. Indeed, for this example, $d^q_W(A, B) = 100$, while $d^p_{MJ}(A, B) = 100\left(\frac{1}{9}\right)^{\frac{1}{p}}$. That is, only the latter semi-metric assigns these remarkably similar sets of structural breaks with a low distance, hence a high affinity.

The Wasserstein metric would grant two time series that have quite similar structural breaks (and hence erratic behavior profiles) low affinity, and hence could include both in a portfolio. This would be a mistake, as such structural breaks as given by A and B in fact have almost all elements in common and should be assigned high affinity, so the portfolio will not choose them both. We illustrate an example of this in Section 4.1.

4. Simulation Study

In this section, we perform two experiments involving simulated time series with specified structural breaks. The first experiment illustrates the computation of the similarity between sets of structural breaks, comparing our MJ_p distance with the Hausdorff and Wasserstein metrics. The primary purpose of this experiment is to illustrate the benefits of our chosen discrepancy measure over alternative and existing metrics. Our examples are chosen to exemplify, via a small number of time series, how the properties proven in Section 3 matter for real data. The second experiment illustrates the allocation of assets in a sample optimization problem, together with constraints typical of an investment policy statement.

4.1. Synthetic Data Simulation

First, we simulate a collection of time series $x_1, \ldots, x_n$ from a GARCH model (Lamoureux and Lastrapes 1990) with m structural breaks determined by jumps at the points $\tau_1, \ldots, \tau_m$. Each time series x_i follows a Student-t distribution with certain specified mean and variance functions. The mean function μ_t contains an autoregressive AR(1) process and a jump component; the latter is a product of a jump direction and magnitude with Bernoulli and gamma distributions, respectively. The variance function σ^2_t contains several terms: an order one short-term component, a long-term persistence component, and a leverage effect component. We display four simulated simulated time series x_i with specified sets of structural breaks τ_j in Figure 1.

Next, we compute the distance matrix and associated affinity matrix between the four synthetic time series, relative to the Hausdorff metric, Wasserstein metric, and $MJ_{0.5}$ and MJ_1 semi-metrics, respectively, and display them in Tables 1–4. These tables collectively illustrate the advantages of the MJ_p semi-metrics compared with the Hausdorff and Wasserstein metrics first discussed in Section 3. First, the Wasserstein metric gives the lowest affinity score between the Time Series 1 and 2, which have 8 out of 9 of their structural breaks in common. These remarkably similar sets of structural breaks are given much lower distance and hence higher affinity under the $MJ_{0.5}$ and MJ_1 distances, illustrating Proposition 4. Next, the Hausdorff metric is far too sensitive to outliers; while TS1 and TS3 have 8 out of 9 points in common, these two time series are given the highest Hausdorff distance among the collection, hence an affinity equal to 0. Similarly, TS2 and TS3 have 7 of 9 points in common, but their assigned affinity is 0.1. The Wasserstein, $MJ_{0.5}$, and MJ_1 distances all recognize the similarity between TS1 and TS3 (as well as TS2 and TS3), with high affinity scores. Once again the $MJ_{0.5}$ and MJ_1 perform better than the Wasserstein in discerning the strong similarity between these time series. We provide the time series TS4 as a reference time series that is quite distinct in its structural breaks from TS1, TS2, and TS3. Only the $MJ_{0.5}$ and MJ_1 assign TS1, TS2, and TS3 mutually high affinity scores, and only an algorithm using them to measure distances between sets of structural breaks would diversify away from including an unsuitably high quantity of TS1, TS2, and TS3 in one asset portfolio. Thus, this simple example of just four synthetic time series highlights the advantages of the MJ_p semi-metric over the Hausdorff and Wasserstein metrics, illustrating the theoretical properties proven in Propositions 3 and 4.
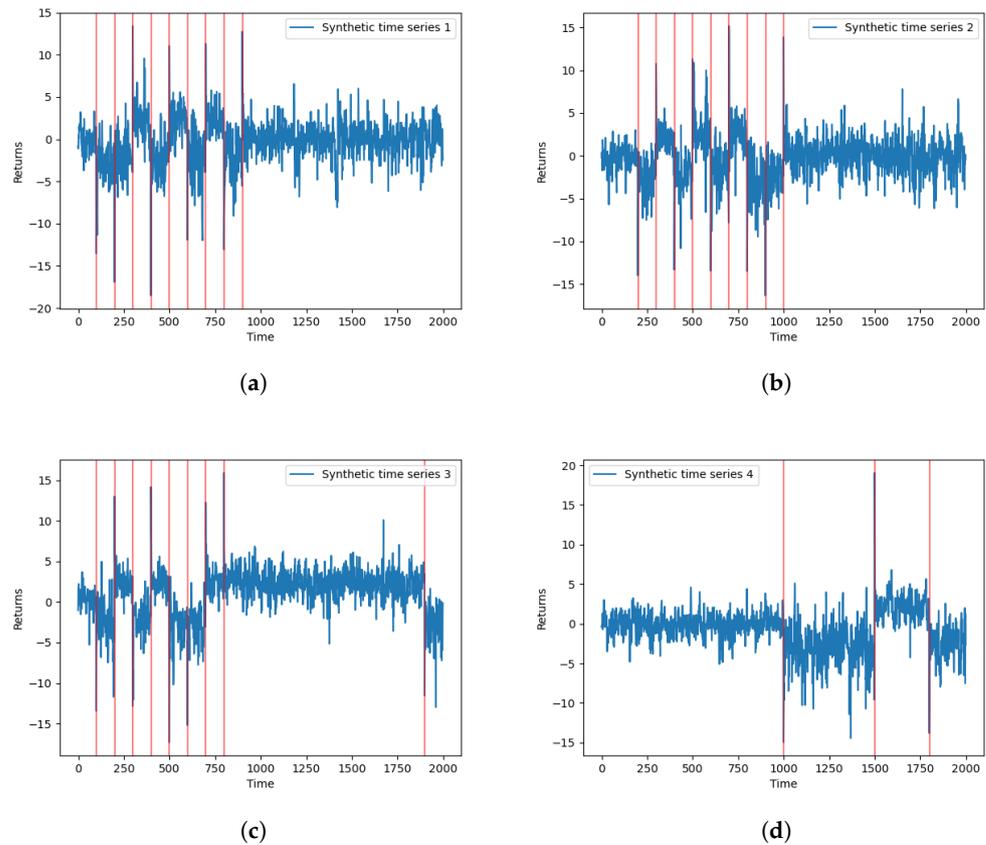
(**a**)   (**b**)

(**c**)   (**d**)

Figure 1. Four synthetic time series (**a**–**d**) exhibiting dependence and correlation in jump behavior. The red annotated lines represent specified structural breaks. The structural breaks are chosen carefully to relate to Corollary 1 and Proposition 4, where we discuss some undesirable properties of the Hausdorff and Wasserstein metrics, respectively. (**a**,**b**) are chosen to have all but one break in common; the Wasserstein metric allocates an excessive distance. (**c**) is chosen to feature a single outlier break compared with (**a**,**b**); the Hausdorff metric is excessively sensitive to the outlier element. As we discuss in Section 4.1 and show in Tables 1–4, only the MJ_p family of semi-metrics properly identifies (**a**–**c**) as highly similar, with (**d**) the one set of structural breaks meaningfully different to the others.

Table 1. Hausdorff distance matrix and affinity matrices between four synthetic time series structural breaks. This allocates an unsuitably excessive distance between time series (**a**,**c**) of Figure 1, illustrating the Hausdorff metric's sensitivity to outliers.

$$
D = \begin{bmatrix} 0 & 100 & 1000 & 900 \\ 100 & 0 & 900 & 800 \\ 1000 & 900 & 0 & 900 \\ 900 & 800 & 900 & 0 \end{bmatrix} ; A = \begin{bmatrix} 1 & 0.9 & 0 & 0.1 \\ 0.9 & 1 & 0.1 & 0.2 \\ 0 & 0.1 & 1 & 0.1 \\ 0.1 & 0.2 & 0.1 & 1 \end{bmatrix}
$$

Table 2. Wasserstein distance matrix and affinity matrices between four synthetic time series structural breaks. This allocates an unsuitably excessive distance between time series (**a**,**b**) of Figure 1, despite their high intersection, showing an undesirable property when measuring discrepancy between sets.

$$
D = \begin{bmatrix} 0 & 100 & 111 & 933 \\ 100 & 0 & 189 & 833 \\ 111 & 189 & 0 & 844 \\ 933 & 833 & 844 & 0 \end{bmatrix} ; A = \begin{bmatrix} 1 & 0.89 & 0.88 & 0 \\ 0.89 & 1 & 0.78 & 0.11 \\ 0.88 & 0.78 & 1 & 0.10 \\ 0 & 0.11 & 0.10 & 1 \end{bmatrix}
$$

Table 3. $MJ_{0.5}$ distance matrix and affinity matrices between four synthetic time series structural breaks. This appropriately identifies (**a**,**b**) as highly similar, (**c**) a little bit further away, and (**d**) the clear outlier element.

$$D = \begin{bmatrix} 0 & 1 & 5 & 461 \\ 1 & 0 & 13 & 306 \\ 5 & 13 & 0 & 327 \\ 461 & 306 & 327 & 0 \end{bmatrix}; A = \begin{bmatrix} 1 & 0.998 & 0.989 & 0 \\ 0.998 & 1 & 0.972 & 0.336 \\ 0.989 & 0.972 & 1 & 0.291 \\ 0 & 0.336 & 0.291 & 1 \end{bmatrix}$$

Table 4. MJ_1 distance matrix and affinity matrices between four synthetic time series' structural breaks. This appropriately identifies (**a**,**b**) as highly similar, (**c**) a little bit further away, and (**d**) the clear outlier element.

$$D = \begin{bmatrix} 0 & 11 & 61 & 517 \\ 11 & 0 & 72 & 417 \\ 61 & 72 & 0 & 367 \\ 517 & 417 & 367 & 0 \end{bmatrix}; A = \begin{bmatrix} 1 & 0.98 & 0.88 & 0 \\ 0.98 & 1 & 0.86 & 0.19 \\ 0.88 & 0.86 & 1 & 0.29 \\ 0 & 0.19 & 0.29 & 1 \end{bmatrix}$$

4.2. Synthetic Data: Portfolio Optimization Experiments

In this section, we apply our portfolio optimization methodology to synthetic data and illustrate the resulting allocation of assets. We generate eight synthetic time series with specified structural breaks. For simplicity, Assets 1–3 have identical numbers of change points with identical locations, as do Assets 4–6; Assets 7 and 8 are outliers. In addition, we set all time series to have the exact same historical return $\mathbb{E}R_i$, so that the numerator of (7) is a positive constant, regardless of the selection of weights. Thus, maximizing the MJ ratio (7) is equivalent to minimizing its denominator $w^T A w$. We display the synthetic time series, together with structural breaks, in Figure 2. We allocate our portfolio subject to a typical constraint $5\% \leq w_i \leq 40\%, \forall i = 1, ..., 8$. Such upper and lower bounds are frequently used in real-world investment policy statements (Coffey 2016), and may be tightened if desired.
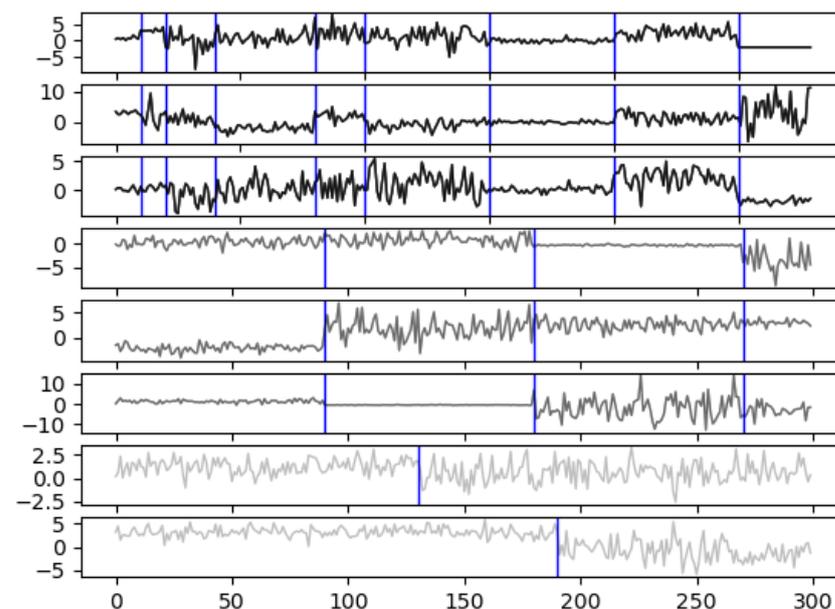

Figure 2. Synthetic time series with fixed historical returns and specified structural breaks, identifying distributional changes. As the historical returns are kept constant, only the sets of structural breaks are used in the objective function (7). These time series form the basis of the synthetic experiment in Section 4.2.

Table 5 records the weights subject to the aforementioned constraints and conditions. The results demonstrate that our optimization framework is able to produce a more even distribution of change points across the portfolio. Assets 7 and 8, with significantly different breaks from the rest of the collection, are allocated more weight: 33.5% and 30.7%, respectively. Much less weight, 6.9%, is allocated to Assets 1, 2, and 3, and just 5% is allocated to Assets 4, 5, and 6. This experiment demonstrates that the algorithm provides diversification with regards to highly correlated structural breaks. Traditional mean variance portfolio optimization would be unable to do so.

Table 5. Portfolio allocation results for synthetic data experiment in Section 4.2. Our method chooses high weights in Assets 7 and 8 to diversify away from the high similarity in structural breaks among Assets 1–6.

Asset	Number of Change Points	Weights
$Asset_1$	8	6.9%
$Asset_2$	8	6.9%
$Asset_3$	8	6.9%
$Asset_4$	3	5%
$Asset_5$	3	5%
$Asset_6$	3	5%
$Asset_7$	1	33.49%
$Asset_8$	1	30.7%

5. Real Data Results

In this section, we apply our methodology to real financial data. We envisage this method being suitable in an asset allocation context, so we use indices and commodities as our underlying candidate investments. We are essentially simulating the role of an asset allocator, such as a pension fund or endowment, interested in macroeconomic asset allocation decisions. There are eight assets we allocate between to illustrate our method: the S&P 500, Dow Jones Index, Nikkei 225 Index, BOVESPA Index, Stoxx 50 Index, ASX 200, oil spot price, and gold spot price, all between January 2009 and November 2019. There are several important details and assumptions in our experiments on real data:

1. We train our algorithm over a relatively long period to estimate the true dynamics between various assets' structural breaks as precisely as possible. Training the algorithm on longer periods provides a more accurate assessment of similarity in varying market dynamics.

2. However, there is a balance between going back far enough to learn appropriate dynamics between asset classes and using too much history that relationships between assets no longer behave the way in which they were estimated. The behavior of individual asset classes and their relationships may change over time.

3. The period from January 2018–June 2019 is a suitable out-of-sample period to test the algorithm, due to the varied market conditions. Most of 2018 provided relatively buoyant equity market returns, with a sharp drop in December 2018, followed by a prolonged recovery until June 2019. We wish to examine how candidate portfolios will perform in various market conditions, particularly in the presence of large drawdowns. In addition, we do not wish to test our algorithm during a period that is too similar to the training interval, as performance could be artificially strong. Thus, this is a suitable period to compare the optimization algorithms' performance.

4. We did not include the COVID-19 market crisis in our test data to ensure that our training data have broadly similar dynamics to the out-of-sample data set. We include a targeted analysis of the COVID-19 crisis in Section 5.3.

5. The role of asset allocation is often guided by an investment policy statement that provides upper and lower bounds for capital allocation decisions. This is captured in the candidate weights' constraints. During pronounced bull and bear markets,

institutional asset allocators may not have the flexibility to implement global optimization solutions. For example, if two asset classes had significantly higher returns and lower volatility than the remainder of candidate investments, the unconstrained solution would allocate all portfolio weight into these two assets. Investment weighting constraints prevent these contrived scenarios from occurring. For our constraints, we place a minimum 5% and maximum 25% of portfolio assets in any candidate investment. This is one of several typical constraints imposed in real-world policy statements—indeed, investment policy statements may include this as their only constraint (Coffey 2016). As mentioned in Section 3, we may impose additional constraints by combining with other optimization methods cited in Section 1.

6. Our method provides an advantage over the simple correlation measure by addressing all three limitations in Section 2. One possible drawback to our proposed method, however, is that to learn meaningful relationships between assets' structural breaks, a long time series history is needed, preferably with many structural breaks observed.

7. When considering portfolio risk in an optimization framework, investors have a variety of measures they may choose to optimize over. Standard deviation, β, downside deviation, and tracking error are just several of these. Our CPO model introduces a mathematical framework that addresses peak-to-trough (drawdown) losses and erratic behavior as a measure of risk. Specifically, the model captures simultaneous asset shocks and aims to minimize the size of drawdowns by creating a uniform spread of change points across all portfolio holdings. We are unaware of any existing measure with these properties.

5.1. Training and Validation Procedure

We train the algorithm between January 2009–December 2017 and test its performance on data from January 2018–June 2019. The training procedure learns the weights allocated to each candidate investment using the aforementioned objective function and constraints. We compare our change point optimization method (CPO) with nine other methods. These methods have been chosen as the best representation of comparative methods in the portfolio optimization literature. Given the breadth of research undertaken within the massive field of econometric portfolio optimization, no comparison with other methods can be completely exhaustive. This list is appropriate, as we judiciously chose methods that cover the most fundamental, best-known, and widely understood objective functions. They also include measures such as conditional value at risk, one of the fundamental measurements of tail risk in the field of econometrics. Our method has value in consciously pursuing an alternative mathematical attribute of diversification, which could also be incorporated with existing methods by portfolio averaging.

First, we apply the Mann–Whitney change point detection algorithm to the training data (log returns between January 2009–December 2017), identifying the locations of structural breaks in the mean for each possible asset. This yields eight sets of change points, where each point is indexed by time. Following Section 2, we apply the $MJ_{0.5}$ semi-metric to determine the distance between candidate sets of breaks. We optimize the MJ ratio objective function in (7) with respect to the weights, determining candidate weight allocations. Finally, we run an out-of-sample forecasting procedure using the weights estimated in our training data. We compare the predictive performance of the ten candidate methodologies between January 2018–June 2019. In addition, we apply agglomerative hierarchical clustering (Müllner 2013) to the resulting distance matrix between assets as exploratory analysis of their similarity with respect to structural breaks.

This identifies a cluster of four highly similar assets (S&P 500, Dow Jones, Stoxx 50, and oil), a cluster of three moderately similar assets (BOVESPA, Nikkei 225, and ASX 200), and an outlier in gold, displayed in Figure 3. These results confirm financial intuition and documented relationships between asset classes, in particular gold's properties as a safe-haven asset. Both the S&P 500 and Dow Jones Index are determined to be in the same

cluster and are accordingly quite similar. Given that there is significant overlap in the constituents of both indices, this is a logical finding.
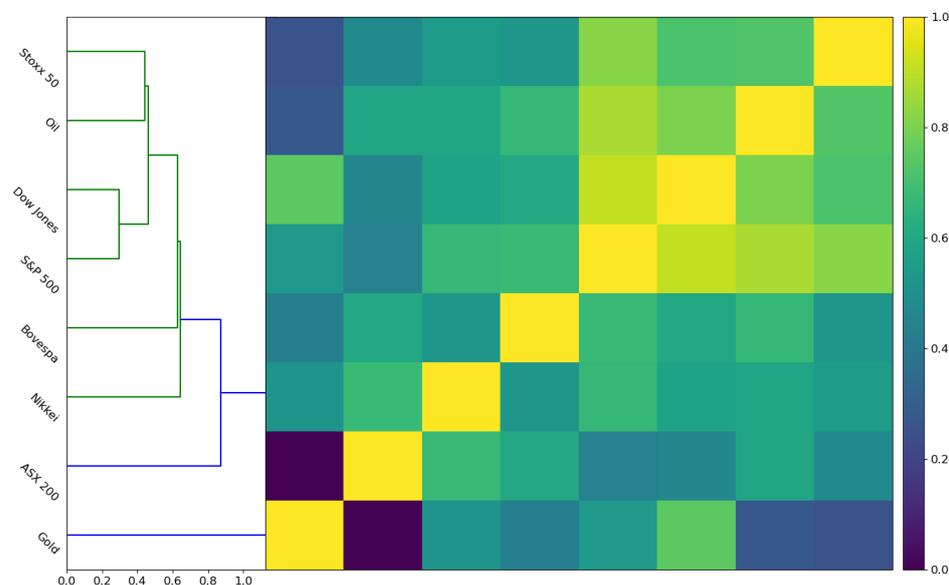

Figure 3. Hierarchical clustering applied to the distance matrix D_{ij} between eight real assets' structural breaks, as defined in Section 2. Results indicate the structural breaks are most similar among the Dow Jones, S&P 500, Stoxx 50, and oil. Gold is the most anomalous asset regarding structural break propagation.

5.2. Out-of-Sample Performance and Distributional Properties

Now, we compare all ten methods' out-of-sample performance, displaying cumulative returns over time in Figure 4 and documenting key metrics in Table 6. With respect to cumulative returns, we see two clear outliers: the entropic value at risk (EVAR) and the conditional value at risk (CVAR) methods, which outperform and underperform, respectively.

Table 6. Results of our change point optimization (CPO) and nine other commonly used portfolio optimization methodologies applied to real data on the test period January 2018–June 2019. Utilized methods are mean-variance optimization (MVO), mean-semivariance (MSV), mean absolute deviation (MAD), first and second lower partial moment (FLPM and SLPM), conditional value at risk (CVAR), entropic value at risk (EVAR), conditional drawdown at risk (CDAR), and Ulcer index (UCI). CPO and EVAR are the two best performing methods. CPO yields the lowest standard deviation, second-lowest drawdown and second-best return. Negative Sharpe ratios convey no useful information, so are omitted. For investors hoping to reduce portfolio volatility, the proposed CPO method is a suitable choice.

Method	Cumulative Returns	Standard Deviation	Sharpe Ratio	Drawdown	Kurtosis
CPO	107.04	**0.0045**	0.99	8.83	1.06
MVO	98.64	0.0060	-	17.08	1.54
MSV	105.76	0.0055	0.66	**6.61**	1.61
MAD	102.28	0.0068	0.21	13.03	2.38
FLPM	101.82	0.0063	0.18	11.05	**0.90**
SLPM	102.26	0.0062	0.23	10.59	0.90
CVAR	72.32	0.0061	-	29.23	0.90
EVAR	**148.55**	0.0053	**5.77**	27.17	1.57
CDAR	100.66	0.0066	0.063	14.22	2.25
UCI	100.60	0.0055	0.069	12.47	1.61

For this purpose, Figure 4a displays just EVAR, CVAR, and our proposed change point optimization method (CPO). We see that CPO produces more stable return trajectories than each outlier method, albeit less total returns than EVAR. Subsequently, we exclude EVAR and CVAR and compare the remaining eight methods in Figure 4b. Among these methods, the CPO method generates the greatest cumulative returns and the lowest standard deviation. For a candidate investor most focused on generating significant returns with minimal volatility, CPO exhibits the most favorable risk–return profile.

In Figure 4c, we contrast the density of daily returns for CPO and three commonly applied portfolio optimization methodologies: mean–variance optimization (MVO), mean–semivariance (MSV), and CVAR. The thinner tails exhibited by CPO show that this method provides consistently reduced volatility in returns. Together with the fact that CPO provides the highest cumulative returns of these methods, we see CPO provides a superior risk-adjusted return compared with these comparable competitive measures. This makes it the most desirable portfolio among all non-outlier candidate methods.

All methods' cumulative returns are documented precisely in Table 6. CPO produces the second-highest cumulative returns (and second-highest Sharpe ratio) when examined on our test data, with a final value of 107.04. The best-performing method is EVAR, which generates a final value of 148.55. Interestingly, the EVAR method also exhibits the highest volatility (standard deviation) and one of the highest drawdowns. This suggests that although CPO did not produce the highest returns, it may have still been the most preferred methodology for investors concerned with volatility.

Next, we examine the out-of-sample standard deviation performance among all comparative methods. CPO produces the lowest standard deviation, with a score of 0.0045. The second-best-performing method is the MSV, with a standard deviation of 0.0055, and the worst-performing methods are the mean absolute deviation method (MAD) and EVAR. That is, despite the outlier method's performance in returns, it exhibits significant volatility.

CPO also performs very strongly with respect to portfolio drawdown. CPO produces the second-lowest drawdown, with a total score of 8.83. The best-performing comparative method is MSV, with a drawdown of 6.61, while CVAR and EVAR have the most significant drawdowns of 29.23 and 27.17, respectively. Again, for investors concerned with strong performance and minimal risk—the CPO and MSV methods produce the most favorable profiles. High levels of drawdown can be particularly concerning for portfolio managers, as respective clients actively tracking their investments may panic during drawdowns and request a withdrawal of funds. For active managers, this is particularly concerning, as there is a heightened chance of funds dropping below the high watermark.

Finally, we turn to kurtosis, where CPO exhibits a level that is approximately average among all the examined methodologies. Compared with other methods that generate strong returns profiles such as EVAR and MSV, CPO has a kurtosis of 1.06, while EVAR and MSV produce scores of 1.57 and 1.61, respectively. Given the positive skew in the CPO distribution, the suppressed kurtosis value is likely indicative of less tail risk in the CPO predictive distribution. The best-performing methods are CVAR and the first and second lower partial moment methods (Omega and Sortino ratios, respectively), all producing scores of 0.9. Although this is indicative of less tail risk, the kurtosis scores of the predictive distribution are likely lower due to reduced average daily returns.
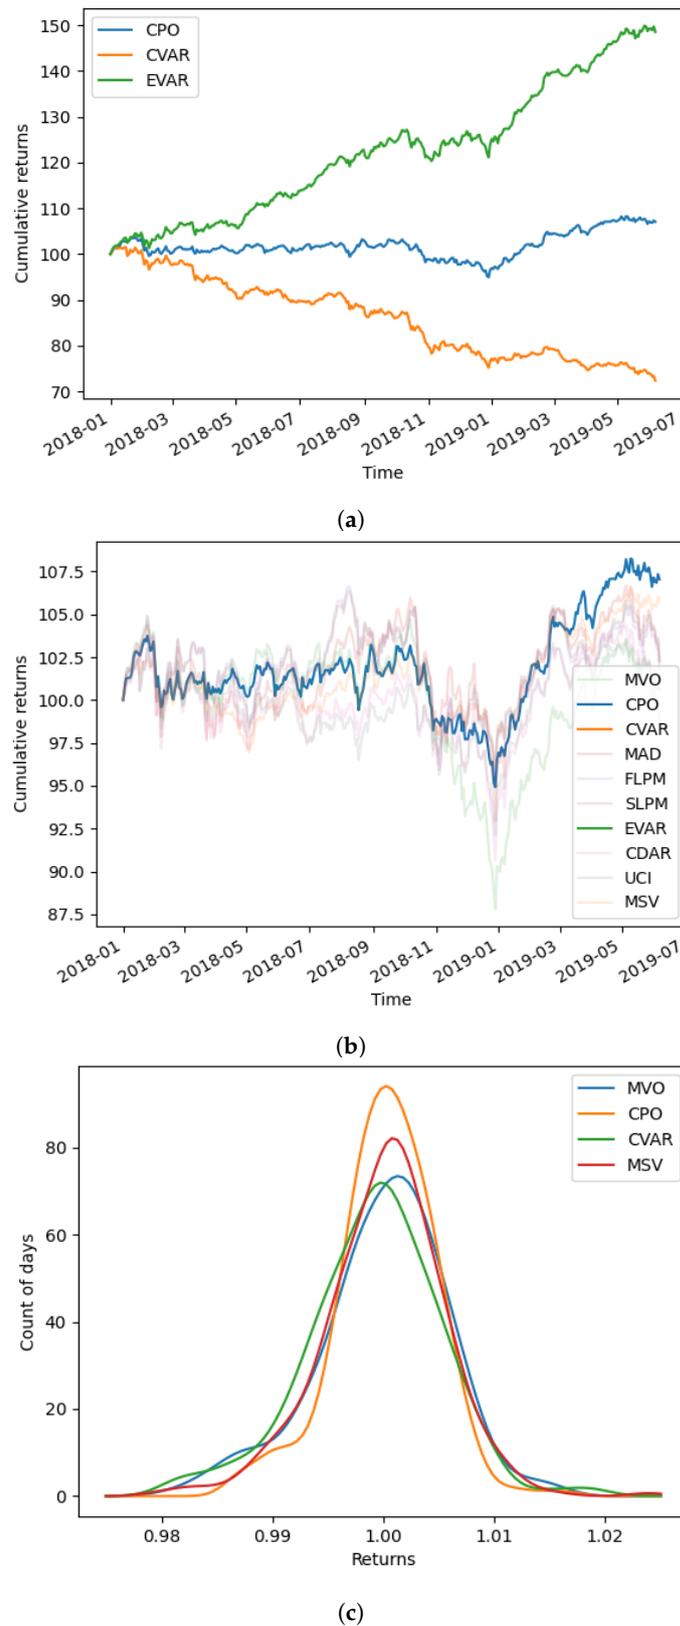
(**a**)

(**b**)

(**c**)

Figure 4. Out-of-sample performance during the January 2018–June 2019 period comparing (**a**) returns of CPO with outliers CVAR and EVAR, (**b**) returns of all methods excluding CVAR and EVAR, and (**c**) distributions of MVO and CPO with commonly used methods MVO, MSV, and CVAR. Our methodology, CPO, generates the second-greatest returns, second-greatest Sharpe ratio, lowest volatility, and second-lowest drawdown of all methods.

5.3. Performance during COVID-19

We devote this section specifically to an analysis of the performance of our methodology (and several others) during the market crash associated with the onset of the COVID-19 pandemic in early 2020 (Akhtaruzzaman et al. 2020; James and Menzies 2021b; Okorie and Lin 2020). We focus on the period of January–June 2020, which was associated with a significant market decline and then correction, including some of the worst one-day market drops in history (Imbert and Franck 2020). We use the same methods as in the previous Section 5.2 and investigate the performance of the corresponding portfolios, renormalized to begin at 100% as of the beginning of January 2020. We display cumulative returns over time in Figure 5 and document the same metrics as before in Table 7.

Similar to Section 5.2, we see again that CPO produces more stable return trajectories than any other method (albeit with CDAR ultimately producing higher returns at the end of the period). CPO produces the second-highest cumulative returns when examined on our COVID-19 data, with a final value of 99.76. Only CDAR produces ultimately higher returns (101.43), at the cost of higher volatility, drawdown, and kurtosis. Again, for a candidate investor focused on weathering a market crisis, CPO exhibits the most favorable risk–return profile. We remark that all optimization methods under investigation share a broadly similar cumulative returns trajectory. This suggests that, regardless of the weight allocations (which vary meaningfully across methods), completely avoiding the market crash is virtually impossible in a long-only setting. One can see that our proposed methodology does an excellent job at mitigating downside risk during a crisis, while maintaining solid positive returns during the subsequent market recovery.

In particular, CPO produces the best (minimal) volatility, drawdown, and kurtosis among all methods, as seen in Table 7. Indeed, CPO carries the lowest standard deviation of 0.018, while the second-lowest standard deviation is 0.025. As in Section 5.2, the suppressed kurtosis of CPO is likely indicative of less tail risk in returns distribution. Combined with the minimal drawdown, this is an especially welcome feature for investors seeking a portfolio to best weather a market crisis, where excessive drawdowns can cause panic and bank runs.
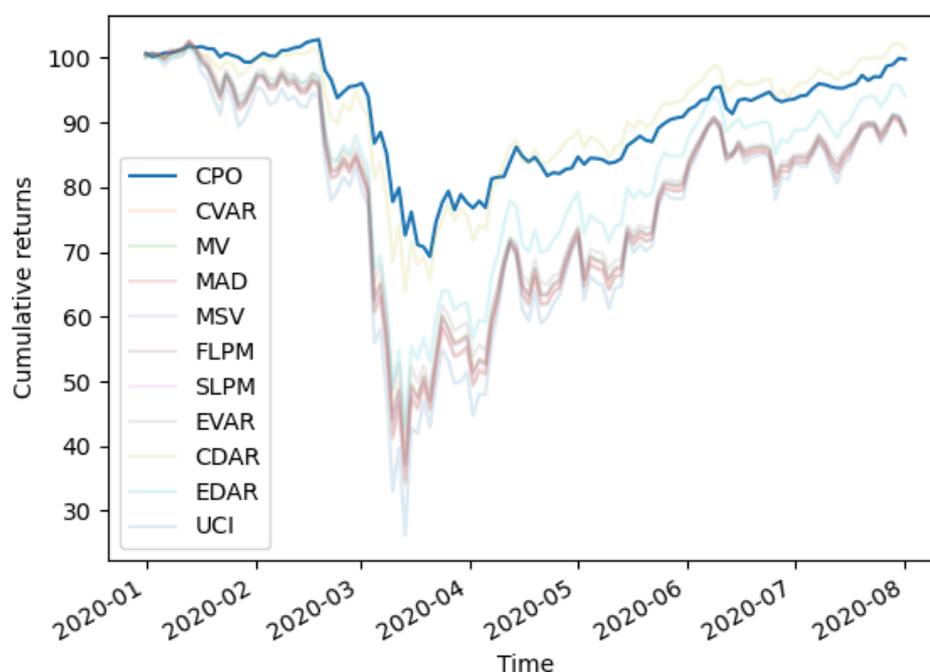

Figure 5. Out-of-sample performance during the height of the COVID-19 market crisis: January–August 2020. We compare the returns of CPO with all methods. Our methodology, CPO, generates the second-greatest returns, lowest volatility, lowest kurtosis, and lowest drawdown of all methods.

Table 7. Results of our change point optimization (CPO) and nine other commonly used portfolio optimization methodologies applied to real data during the height of the COVID-19 market crisis: January–August 2020. Utilized methods are mean–variance optimization (MVO), mean–semivariance (MSV), mean absolute deviation (MAD), first and second lower partial moment (FLPM and SLPM), conditional value at risk (CVAR), entropic value at risk (EVAR), conditional drawdown at risk (CDAR), and Ulcer index (UCI). CPO and CDAR are the two best-performing methods. CPO yields the lowest volatility, drawdown, and kurtosis, and the second best return. For investors hoping to reduce portfolio volatility, the proposed CPO method is a suitable choice. As almost all returns are negative, the Sharpe ratio is uninformative and has been omitted.

Method	Cumulative Returns	Standard Deviation	Drawdown	Kurtosis
CPO	99.76	**0.018**	**33.55**	**8.55**
MVO	88.82	0.036	65.13	15.48
MSV	88.69	0.036	65.54	15.60
MAD	88.06	0.038	67.71	15.99
FLPM	88.08	0.038	68.39	16.11
SLPM	88.69	0.036	65.51	15.60
CVAR	88.45	0.035	64.41	15.54
EVAR	89.12	0.033	60.99	14.90
CDAR	**101.43**	0.025	38.38	9.49
UCI	88.40	0.045	75.58	17.15

5.4. Sampling Study of Structural Breaks between Countries' Financial Indices

In this section, we conduct a sampling study to investigate patterns in the collective distance between financial assets. For robustness, we draw data from a completely independent selection of assets as the prior sections, analyzing log returns data of the national financial indices of 19 countries. Our chosen countries are Australia, Brazil, Canada, China, France, Germany, India, Indonesia, Italy, Japan, Korea, the Netherlands, Russia, Saudi Arabia, Spain, Switzerland, Turkey, the United Kingdom (UK), and the United States (US), with data ranging from 2001–2020. We begin by applying our change point algorithm (described in Appendix A) to obtain sets of structural breaks $S_i, i = 1, \ldots, 19$ for each country.

Next, we perform a variety of repeated sample experiments. We vary $n = 4, 6, 8, 10, 12, 14, 16$ and for each value of n, sample $K = 2000$ draws of n countries (without replacement) from the collection. A draw of n countries produces an $n \times n$ distance matrix D (using the $\mathrm{MJ}_{0.5}$ semi-metric as elsewhere in the paper) between the countries. We compute the following normalized L^1 norm of the matrix D to measure the collective magnitude of all distances between countries:

$$\|D\| = \frac{1}{n(n-1)} \sum_{i,j=1}^{n} |D_{ij}|. \tag{34}$$

We normalize by the number of nonzero elements in this distance matrix, $n(n-1)$. Thus, the sampling procedure produces $K = 2000$ values of $\|D\|$ for each size n. In Figure 6, we show the distributions of values for each n. We also record a 90% interval consisting of the 5th and 95th quantile in Table 8. As n increases, we observe an increase in the overall mean of the distribution—this effect is relatively rapid at first but slows after $n = 8$. We also observe a relatively quick increase in the lower limit and a slower decrease in the upper limit. This is to be somewhat expected. In small samples (such as $n = 4$), it is possible to repeatedly by chance select just four countries that are relatively similar to each other in terms of structural breaks. However, selecting 8 or 10 countries that are all similar to each other (and hence yielding a small value of $\|D\|$) is much less likely, producing a greater value of the lower limit.
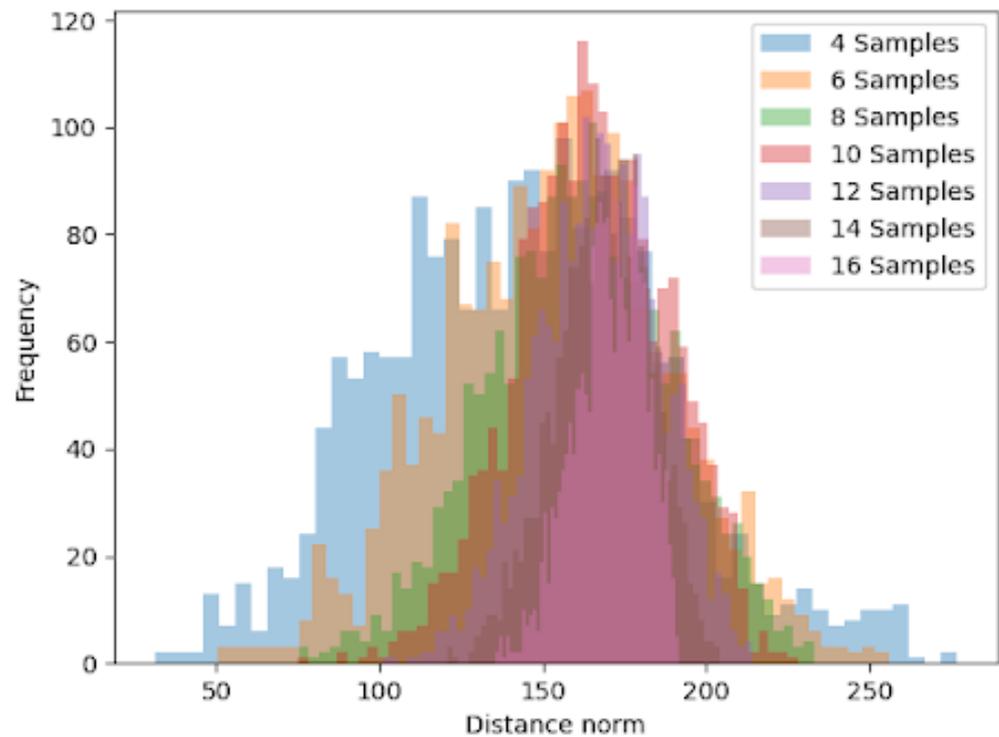
Figure 6. Distributions for the normalized distance matrix norm $\|D\|$ when repeatedly drawing $K = 2000$ samples of size n from our collection of 19 countries and computing the distances between structural breaks. The range of values increases but tightens in variance as n increases, reflecting that the normal benefits of diversification when increasing portfolio size also apply to the quantitative measure of distances between structural breaks.

Table 8. Ninety percent confidence intervals for the normalized distance matrix norm $\|D\|$ when repeatedly drawing $K = 2000$ samples of size n from our collection of 19 countries and computing the distances between structural breaks. The range of values increases but tightens in variance as n increases, reflecting that the normal benefits of diversification when increasing portfolio size also apply to the quantitative measure of distances between structural breaks.

Sample Size	Lower Limit	Upper Limit
4	72.99	213.17
6	100.70	210.92
8	113.35	208.32
10	127.02	201.08
12	135.86	196.36
14	143.81	191.59
16	151.42	186.31

The next key finding in this experiment is the significant reduction in the spread of the distribution of matrix norms, as a larger number of countries are sampled. The light blue coloring, which is a relatively diffuse distribution, corresponds to when only four country indices are sequentially sampled. This distribution is centered ∼150, with a total range of approximately 250. As the number of stocks sampled increases, we see the distribution of norm values become successively narrower (exhibiting less variance). This culminates in the pink distribution ($n = 16$), centered around ∼175, with a total range of approximately 50.

6. Discussion

We proposed a novel optimization method, which utilizes semi-metrics between sets of structural breaks, to reduce simultaneous asset shocks across an investment portfolio. This should be understood as an alternative approach to avoid simultaneous losses that can gravely threaten an investor's holdings. Experiments on synthetic data confirm that we are able to detect similar time series in terms of structural breaks and accordingly allocate highly similar assets less portfolio weight. In addition, synthetic experiments illustrate concretely the proven theoretical properties of our particular choice of semi-metric and its benefits over other existing metrics between finite sets. Experiments on real data suggest that our method may significantly reduce both portfolio volatility and drawdown when compared with numerous existing methodologies.

This novel optimization framework may have significant implications for asset allocation and portfolio management professionals who are interested in alternative measures of risk. Our method diversifies well away from portfolio drawdown and seeks to avoid the erratic behavior of highly clustered change points. Our method is flexible, and different change point algorithms may be married with other distance measures or objective functions for alternative approaches. Finally, our method's efficient implementation and reasonable complexity mean it easily generalizes to large instances of the portfolio optimization problem, especially when convex real-world constraints are applied. In the event of more difficult constraints, the calculation of our distances between sets of change points could be combined with efficient optimization approaches in the econometrics literature.

There are several limitations in our optimization framework. Change point detection methodologies vary widely, and there is substantial literature on their advantages and disadvantages (Gustafsson 2001; Hawkins and Zamba 2005). One potential limitation of our chosen change point detection algorithm is its deterministic selection (via hypothesis testing and maximization of test parameters) of change point locations. Alternative approaches, such as Bayesian methods, may provide a probabilistic approach incorporating the uncertainty around change points' existence. Next, our methodology requires a long training period to learn meaningful relationships between assets' structural breaks, and it is conceivable that such relationships from asset histories no longer hold in the present. Furthermore, any distance measure between finite sets will have its limitations, such as the Hausdorff metric's sensitivity to outliers, the semi-metrics' failure in the triangle inequality, or the potentially excessive averaging in the MJ_p family. To ameliorate these limitations, we believe model averaging with other methods that require a smaller training time and other mathematical quantities other than distances between finite sets could be beneficial. Furthermore, it is our hypothesis that structural breaks are more likely to have an underlying persistence that is more consistent than their returns. Each asset's structural break propagation most likely has a strong link to the asset's existence within the very complex system of the global economy. Structural change in this manner is likely to occur at a much lower frequency than drivers which dictate changes in market returns and volatility (such as market sentiment).

Future research could work to ameliorate such limitations by building on the proposed framework. Different change point detection algorithms could be used for other stochastic properties, such as the variance in the returns, or to reflect uncertainty in the breaks' existence. One could explore how results change when the order of p changes within the MJ_p semi-metric, or when entirely different distance measures are used between sets. One could conceivably calibrate the value of p, such as selecting a portfolio based on the optimal MJ ratio for various ratios of p, and then using another optimization function such as the Sharpe ratio to tune p as a hyperparameter. More broadly, it is conceivable that one could diversify between sets through other means than reducing their discrepancy to a single scalar value. Even with the judicious choice of p, the outlier sensitivity in the denominator max D of (6) is a limitation of our framework, acknowledged in Remark 1, so alternative constructions other than this affinity matrix could be explored. Furthermore, one could

combine our methodology, which requires a long history of training data, with alternative methods that give more weight to the recent behavior of a financial asset.

This work not only has broad theoretical value, offering for the first time how structural breaks may be used in portfolio optimization and proving numerous properties of our implementation relative to other discrepancy measures between structural breaks, but it also has experimental value beyond our specific application. Model stacking and averaging, also known as ensemble learning (Sagi and Rokach 2018), has proven to have great utility in portfolio optimization (Shen et al. 2019; Ünlü and Xanthopoulos 2021), even combining rather unrelated methods. As we are unaware of portfolio techniques that study structural breaks, we believe model stacking with other methods (focused on variance, tail risk, conditional correlations, and others) could benefit by bringing in another mathematical property of financial time series to assist in diversification and econometric-derived reduction in risk. Alternatively, our method could be combined with numerous methods to clean noise in financial time series prior to analysis. Indeed, for the standard Markowitz model of covariance and correlations, numerous authors have argued that correlation matrices approximate random matrices and should be cleaned (Laloux et al. 1999; Wątorek et al. 2021), with numerous approaches proposed for this purpose (Bun et al. 2017). Conceivably, such preprocessing could be applied in our setting to the sets of structural breaks of different assets.

7. Conclusions

We proposed a new concept of using distances between structural breaks of time series for portfolio optimization and provided a specific implementation. Our first implementation has promising results in its own right and also offers numerous directions for further research by incorporating other advances in the literature, such as noise cleaning and model stacking. More broadly, we hope this paper will invite other novel approaches and concepts towards exploring different "root causes" of simultaneous shocks in financial holdings. Avoiding these has value not just for investors and portfolio managers but for the health of the broader economy as a whole.

Author Contributions: N.J. and M.M. are equal first authors, playing an equal role in every aspect of the manuscript. J.C. performed simulation experiments and provided edits to the manuscript. All authors have read and agreed to the published version of the manuscript.

Abbreviations

The following abbreviations are used in this manuscript:

CPO	Change point optimization method
MVO	Mean–variance optimization
MSV	Mean–semivariance
MAD	Mean absolute deviation
FLPM	First lower partial moment
SLPM	Second lower partial moment
CVAR	Conditional value at risk

EVAR	Entropic value at risk
CDAR	Conditional drawdown at risk
UCI	Ulcer index
CPM	Change point model

Appendix A. Change Point Detection Algorithm

In this section, we describe the change point detection algorithm used in Section 2. The general change point detection framework is the following: a sequence of observations $x_1, x_2, \ldots, x_n$ are drawn from random variables $X_1, X_2, \ldots, X_n$ and undergo an unknown number of changes in distribution at points $\tau_1, \ldots, \tau_m$. We assume observations are independent and identically distributed between change points, that is, between each change point, a random sampling of the distribution is occurring. Ross (2015) notates this as follows:

$$X_i \sim \begin{cases} F_0 \text{ if } i \leq \tau_1 \\ F_1 \text{ if } \tau_1 < i \leq \tau_2 \\ F_2 \text{ if } \tau_2 < i \leq \tau_3, \\ \ldots \end{cases} \tag{A1}$$

While this requirement of independence may appear restrictive, dependence can generally be accounted by several means, such as modeling the underlying dynamics or drift process and then applying a change point algorithm to the model residuals or one-step-ahead prediction errors, as described by Gustafsson (2001). The change point models described below that we apply in this paper follow Ross (2015).

Appendix A.1. Batch Detection (Phase I)

This first phase of change point detection is retrospective. We are given a finite sequence of observations $x_1, \ldots, x_n$ from random variables $X_1, \ldots, X_n$. For simplicity, we assume at most one change point exists. If a change point exists at time k, this means observations have a distribution of F_0 prior to the change point and a distribution of F_1 proceeding the change point, where $F_0 \neq F_1$. Then, one must test between the following two hypotheses for each k:

$$H_0 : X_i \sim F_0, i = 1, \ldots, n \tag{A2}$$

$$H_1 : X_i \sim \begin{cases} F_0 & i = 1, 2, \ldots, k \\ F_1, & i = k+1, k+2, \ldots n \end{cases} \tag{A3}$$

and select the most suitable k.

One proceeds with a two-sample hypothesis test, where the choice of test depends on the assumptions about the underlying distributions. Nonparametric tests can be chosen to avoid distributional assumptions. One appropriately chooses a two-sample test statistic $D_{k,n}$ and a threshold $h_{k,n}$. If $D_{k,n} > h_{k,n}$, then the null hypothesis is rejected and one provisionally assumes that a change point has occurred after x_k. These test statistics $D_{k,n}$ are normalized to have mean 0 and variance 1 and are evaluated at all values $1 < k < n$; the largest value is assumed to be coincident with the existence of our sole change point. The test statistic is then

$$D_n = \max_{k=2,\ldots,n-1} D_{k,n} = \max_{k=2,\ldots,n-1} \left| \frac{\tilde{D}_{k,n} - \mu_{\tilde{D}_{k,n}}}{\sigma_{\tilde{D}_{k,n}}} \right| \tag{A4}$$

where $\tilde{D}_{k,n}$ are non-normalized statistics.

The null hypothesis of no change is rejected if $D_n > h_n$ for an appropriately chosen threshold h_n. In this case, we conclude that a (unique) change point has occurred, and its location is the value of k which maximizes $D_{k,n}$. That is,

$$\hat{\tau} = \underset{k}{\text{argmax}}\, D_{k,n}. \tag{A5}$$

This threshold h_n is chosen to bound the Type 1 error rate, as is commonplace in statistical hypothesis testing. First, one specifies an acceptable level α for the proportion of false positives, that is, the probability of falsely declaring that a change has occurred when in fact it has not. Then, h_n is chosen as the upper α quantile of the distribution of D_n under the null hypothesis. For the details of computation of this distribution, one can see Ross (2015).

The computational cost of this first phase is $O(n)$, where n is the number of observations. Indeed, this calculates and compares $n - 2$ values of the (normalized) test statistics $D_{k,n}, k = 2, \ldots, n - 1$. This is implemented efficiently in C++.

Appendix A.2. Sequential Detection (Phase II)

In this second phase, the sequence $(x_t)_{t \geq 1}$ does not have a fixed length. New observations are continually received over time, and multiple change points may be present. Assuming no change point exists so far, this approach treats $x_1, \ldots, x_t$ as a fixed-length sequence and computes D_t as described in phase I. A change is flagged if $D_t > h_t$ for an appropriately chosen threshold. If no change is detected, the next observation x_{t+1} is brought into the sequence of consideration. If a change is detected, the process restarts from the data point immediately following the detected change point. Thus, the procedure consists of a repeated sequence of hypothesis tests.

In this sequential setting, h_t is selected so that the probability of incurring a Type 1 error is constant over time, so that under the null hypothesis of no change, the following holds:

$$P(D_1 > h_1) = \alpha, \tag{A6}$$

$$P(D_t > h_t | D_{t-1} \leq h_{t-1}, \ldots, D_1 \leq h_1) = \alpha,\ t > 1. \tag{A7}$$

In this case, assuming that no change occurs, the expected number of observations received before a false positive detection occurs is equal to $\frac{1}{\alpha}$. This quantity is often referred to as the average run length, or ARL$_0$. Additional details on appropriate values of h_t are detailed by Ross (2015).

In the context of our paper, this algorithm is performed on time series of length T. Thus, this second phase involves up to T implementations of Phase I, with a new observation x_{t+1} brought in each step. As the complexity cost of Phase I is up to $O(T)$, this means the total cost of the CPM algorithm is $O(T^2)$. This is implemented efficiently in C++.

Appendix B. Overview and Properties of Distances between Sets

In this section, we provide an overview of (semi-)metric distances, with a focus on distance between (finite) sets, and motivate the choice of distance between sets of structural breaks chosen in (4) and (5).

Appendix B.1. Overview of Metrics

We first recall the definition of a metric d on a set X. A pairing $d : X \times X \to \mathbb{R}$ is called a metric if it satisfies the following axioms for all $x, y, z \in X$:

1. $d(x, y) \geq 0$, with equality if and only if $x = y$;
2. $d(x, y) = d(y, x)$;
3. $d(x, z) \leq d(x, y) + d(y, z)$.

d is a *semi-metric* if it satisfies (i) and (ii) but not necessarily (iii), which is known as the *triangle inequality*. If d is a metric on X, then the pair (X, d) is called a metric space (Rudin 1976).

As discussed in Section 1.4, the focus of our methodology is measuring discrepancy between finite sets. With this in mind, the relevant class of (semi-)metrics for this paper is that between subsets of a given metric space. We begin with more explanation of a concept used in Section 2. Let S be a subset of a metric space X, and $x \in X$. Then, the distance from the element x to the set S is defined as the minimal distance from x to any point in S, computed as follows:

$$d(x, S) = \inf_{s \in S} d(x, s). \tag{A8}$$

Now $d(x, S) \geq 0$ with equality if and only if x lies in the closure of S. In addition, $d(-, S); X \to \mathbb{R}$ is continuous. This quantity $d(x, S)$ is the base ingredient of several existing and recently introduced (semi-)metrics between sets.

Appendix B.2. Distances between Sets

Now, let $S, T \subset X$ be (finite) subsets of any metric space. A common first notion of distance between these subsets is defined as the minimal distance between these subsets, defined by

$$d_{\min}(S, T) = \inf_{s \in S} d(s, T) = \inf_{s \in S} \inf_{t \in T} d(s, t) = \inf_{s \in S, t \in T} d(s, t). \tag{A9}$$

Note $d_{\min}(S, T) = 0$ if S, T intersect. In fact, $d_{\min}(S, T) = 0$ if and only if their closures (in the ambient space X) intersect. So, this is not an effective metric between subsets, as it can frequently be zero for sets that are markedly different. We proceed to outline some existing (semi-)metrics between finite sets that have been used for various applications (Conci and Kubrusly 2017).

We begin with the Hausdorff distance, already defined in Definition 1:

$$d_H(S, T) = \max \left(\sup_{s \in S} d(s, T), \sup_{t \in T} d(t, S) \right) \tag{A10}$$

$$= \sup \{ d(s, T), s \in S; d(t, S), t \in T \}. \tag{A11}$$

Essentially, the Hausdorff metric considers how separated S and T are at the most, rather than at least, compared with (A9). More precisely, it is the supremum or L^∞ norm of all minimal distances from points $s \in S$ to T and points $t \in T$ to S, as defined in (A8). The Hausdorff distance satisfies the triangle inequality (so it is a true metric, rather than just a semi-metric), but this supremum is highly sensitive to even a single outlier. Indeed, this is the content of Proposition 3 and Corollary 1, that just one element of S or T can result in great changes to $d_H(S, T)$.

Next, we discuss the pre-existing modified Hausdorff distances, which are semi-metrics that were used in computer vision and other tasks.

Definition A1 (Modified Hausdorff distance 1). *The first modified Hausdorff distance MH_1 is defined as follows (Deza and Deza 2013; Dubuisson and Jain 1994):*

$$d_1^{MH}(S, T) = \max \left(\frac{1}{|S|} \sum_{s \in S} d(s, T), \frac{1}{|T|} \sum_{t \in T} d(t, S) \right). \tag{A12}$$

It takes a first step at replacing the max in the Hausdorff distance with geometric averaging.

Definition A2 (Modified Hausdorff distance 2). *The second modified Hausdorff distance MH_2 is defined as follows (Dubuisson and Jain 1994; Eiter and Mannila 1997):*

$$d_2^{MH}(S, T) = \sum_{s \in S} d(s, T) + \sum_{t \in T} d(t, S). \tag{A13}$$

Unlike (A12), this captures the total deviation between one set and another, with no averaging.

Definition A3 (Modified Hausdorff distance 3). *The third modified Hausdorff distance MH_3 is defined as follows (Deza and Deza 2013; Dubuisson and Jain 1994):*

$$d_3^{MH}(S,T) = \frac{1}{|S|+|T|}\left(\sum_{s\in S} d(s,T) + \sum_{t\in T} d(t,S)\right). \tag{A14}$$

This is a variant of (A12) with a different averaging component, referred to as geometric mean error between two images.

In addition to the Hausdorff metric and the three pre-existing modified Hausdorff distances defined above, there is also the Wasserstein distance, discussed in Section 3.

Now, we turn to the more recently introduced family of semi-metrics introduced in James et al. (2020) and motivate it as our choice for this paper. Here, we first introduced the MJ_1 semi-metric:

$$d_{MJ}^1(S,T) = \frac{1}{2}\left(\frac{\sum_{t\in T} d(t,S)}{|T|} + \frac{\sum_{s\in S} d(s,T)}{|S|}\right). \tag{A15}$$

Our initial motivation for this is as follow: In Dubuisson and Jain (1994), the authors asserted that their distance MH_1 is the best for image matching. To reach this conclusion, they took two steps. First, they compared three favorable operators, f_2, f_3, and f_4, each operating on minimal distances $d(s,T), d(t,S)$, as defined in Equation (A8). They briefly argued that f_2, equivalent to taking the max in the MH_1, is preferable to other operators, citing a "larger spread'.' Second, they argued that a process of averaging distances is superior to taking Kth ranked distances, such as the median. We differed with the first step of their reasoning and replaced the max in their MH_1 with the L^1 norm average of all the minimum distances from S to T and T to S, as seen in (A15). We proceeded to detail some reasons why we preferred the MJ_1 over the three aforementioned modified Hausdorff distances, explained in Propositions 3.1, 3.2, 3.5, and 3.6 of James et al. (2020).

Regarding the second step of their reasoning (Dubuisson and Jain 1994), we agreed that an averaging process was more suitable regarding outlier error than the alternative processes and chose to generalize this by using other L^p norm averages. Thus, we introduced the family of MJ_p semi-metrics. We define the MJ_p distance by

$$d_{MJ}^p(S,T) = \left(\frac{\sum_{t\in T} d(t,S)^p}{2|T|} + \frac{\sum_{s\in S} d(s,T)^p}{2|S|}\right)^{\frac{1}{p}}. \tag{A16}$$

The normalization within the expression is chosen such that

$$d_{MJ}^p(S,T) \le d_H(S,T) \text{ for all } p, \text{ and } \lim_{p\to\infty} d_{MJ}^p(S,T) = d_H(S,T). \tag{A17}$$

Thus, d_H can now be viewed as the L^∞ norm of these distances, that is, our family of semi-metrics includes the Hausdorff distance as a limiting case when $p\to\infty$. So, the existing Hausdorff metric was thus placed in our newly introduced family of semi-metrics. The parameter p sets up a trade-off of sorts: as p gets larger, d_{MJ}^p becomes closer to a metric satisfying the triangle inequality. However, as p gets smaller, Proposition 3 and Corollary 1 show that d_{MJ}^p is less affected by outlier elements.

It is with this prior literature and body of work in mind that, when we wish to measure discrepancy between time series' sets of structural breaks (which are finite sets), the MJ_p semi-metric in (4) was a natural choice. As for the precise selection of p, our optimization framework did not particularly rely on the triangle inequality, while outlier sensitivity is much more important. Hence, we select a small value of p, in this case $p = \frac{1}{2}$.

Appendix B.3. Illustration Study of Different (Semi-)Metrics

In this section, we generate some figures to graphically illustrate different values the aforementioned (semi-)metrics between sets can take. First, we generate a collection of ten time series each, inspired by the synthetic time series in Section 4.1 and experiments in James et al. (2020). The collection is displayed in Figure A1, chosen to feature time series with moderate outlier elements. In this scenario, we consider the first five time series (TS1–TS5 inclusive) as similar, the next three (TS6–TS8) as similar, and the final two (TS9 and TS10) as dissimilar to all other time series.

The graphical representation of these distances is supplied in Figure A2, in which we apply hierarchical clustering to the collection of synthetic time series, using the Hausdorff, MJ_1, $MJ_{0.5}$, and MJ_2 distances. Even in this instance of moderate outliers, the Hausdorff distance (Figure A2a) fails to correctly identify the general structure in the time series collection (one cluster of 1–5, one of 6–8, and two outliers). The remaining three semi-metrics correctly identify the general structure in the time series collection. As predicted, the $MJ_{0.5}$ (Figure A2b) does the best job, both at distinguishing between the two clusters of similarity as well as highlighting the fact that the two outliers are distinct from everything else. We thus provide this as graphical evidence of the suitability of $p = \frac{1}{2}$ to handle outlier elements, a necessary part of our overall optimization framework, as mentioned first in Remark 1.
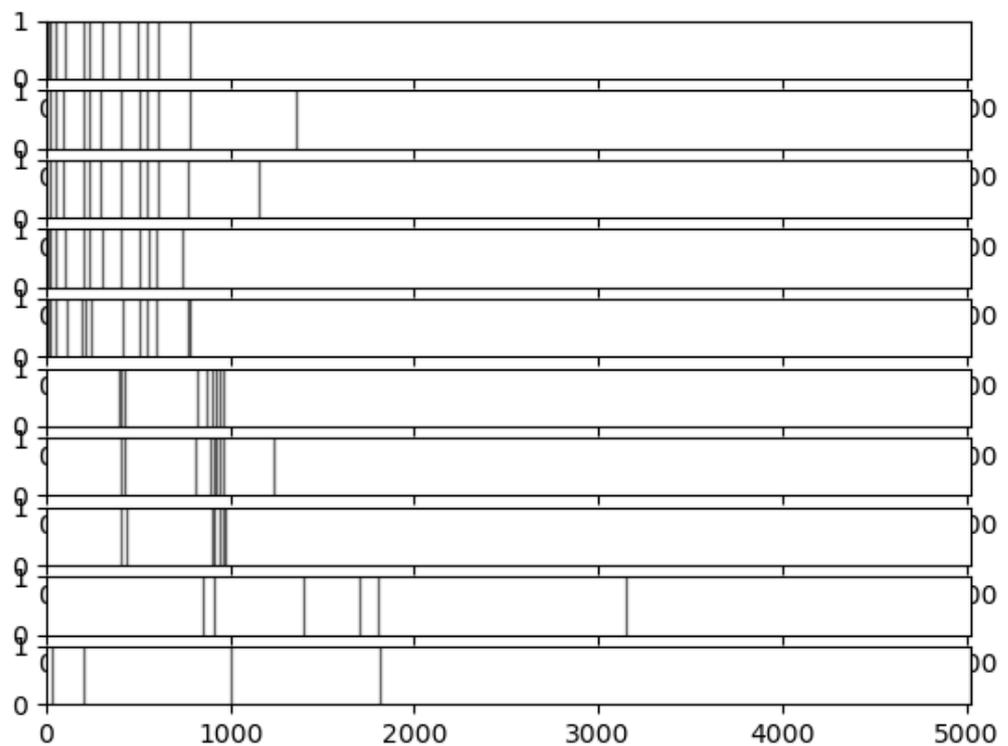

Figure A1. Collection of ten synthetic time series with structural breaks displayed. Two clusters of similarity (TS1–TS5 and TS6–TS8) are observed, as well as two outlier elements (TS9 and TS10) not similar to anything else.
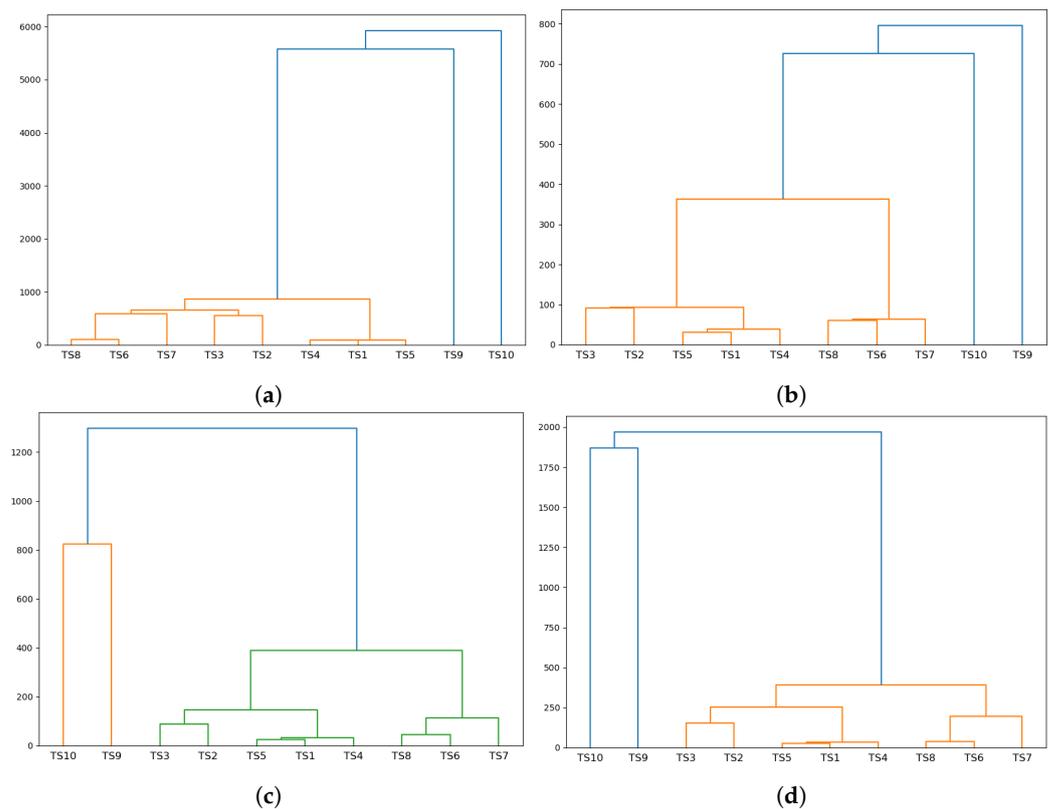
Figure A2. Hierarchical clustering applied to the ten synthetic time series in Figure A1 using (**a**) the Hausdorff metric, (**b**) the $MJ_{0.5}$, (**c**) the MJ_1, and (**d**) the MJ_2. Results indicate that the $MJ_{0.5}$ does the best job at distinguishing between the two clusters of similarity and highlighting the dissimilarity of the two outliers.

References

Adams, Ryan Prescott, and David J. C. MacKay. 2007. Bayesian online changepoint detection. *arXiv*, arXiv:0710.3742.

Akhtaruzzaman, Md, Sabri Boubaker, and Ahmet Sensoy. 2020. Financial contagion during COVID–19 crisis. *Finance Research Letters* 38: 101604. [CrossRef] [PubMed]

Akoglu, Leman, Hanghang Tong, and Danai Koutra. 2014. Graph based anomaly detection and description: A survey. *Data Mining and Knowledge Discovery* 29: 626–88. [CrossRef]

Alexander, Gordon J., and Alexandre M. Baptista. 2002. Economic implications of using a mean-VaR model for portfolio selection: A comparison with mean-variance analysis. *Journal of Economic Dynamics and Control* 26: 1159–93. [CrossRef]

Almahdi, Saud, and Steve Y. Yang. 2017. An adaptive portfolio trading system: A risk-return portfolio optimization using recurrent reinforcement learning with expected maximum drawdown. *Expert Systems with Applications* 87: 267–79. [CrossRef]

Alves, Luiz G. A., Higor Y. D. Sigaki, Matjaž Perc, and Haroldo V. Ribeiro. 2020. Collective dynamics of stock market efficiency. *Scientific Reports* 10: 21992. [CrossRef] [PubMed]

Ammar, E., and H. A. Khalifa. 2003. Fuzzy portfolio optimization a quadratic programming approach. *Chaos, Solitons & Fractals* 18: 1045–54. [CrossRef]

Anagnostopoulos, K. P., and G. Mamanis. 2011. The mean–variance cardinality constrained portfolio optimization problem: An experimental evaluation of five multiobjective evolutionary algorithms. *Expert Systems with Applications* 38: 14208–17. [CrossRef]

Atallah, Mikhail J. 1983. A linear time algorithm for the Hausdorff distance between convex polygons. *Information Processing Letters* 17: 207–9. [CrossRef]

Atallah, Mikhail J., Celso C. Ribeiro, and Sergio Lifschitz. 1991. Computing some distance functions between polygons. *Pattern Recognition* 24: 775–81. [CrossRef]

Baddeley, A. J. 1992. Errors in binary images and an L^p version of the Hausdorff metric. *Nieuw Archief voor Wiskunde* 10: 157–83.

Ballestero, Enrique. 2005. Mean-semivariance efficient frontier: A downside risk model for portfolio selection. *Applied Mathematical Finance* 12: 1–15. [CrossRef]

Barry, Daniel, and J. A. Hartigan. 1993. A Bayesian analysis for change point problems. *Journal of the American Statistical Association* 88: 309. [CrossRef]

Basalto, Nicolas, Roberto Bellotti, Francesco De Carlo, Paolo Facchi, Ester Pantaleo, and Saverio Pascazio. 2007. Hausdorff clustering of financial time series. *Physica A: Statistical Mechanics and Its Applications* 379: 635–44. [CrossRef]

Basalto, Nicolas, Roberto Bellotti, Francesco De Carlo, Paolo Facchi, Ester Pantaleo, and Saverio Pascazio. 2008. Hausdorff clustering. *Physical Review E* 78: 046112. [CrossRef]

Bhansali, Vineer. 2007. Putting economics (back) into quantitative models. *The Journal of Portfolio Management* 33: 63–76. [CrossRef]

Boasson, Vigdis, Emil Boasson, and Zhao Zhou. 2011. Portfolio optimization in a mean-semivariance framework. *Investment Management and Financial Innovations* 8: 58–68.

Bongini, L., M. Degli Esposti, C. Giardinà, and A. Schianchi. 2002. Portfolio optimization with short-selling and spin-glass. *The European Physical Journal B - Condensed Matter* 27: 263–72. [CrossRef]

Braione, Manuela, and Nicolas Scholtes. 2016. Forecasting value-at-risk under different distributional assumptions. *Econometrics* 4: 3. [CrossRef]

Brass, Peter. 2002. On the nonexistence of Hausdorff-like metrics for fuzzy sets. *Pattern Recognition Letters* 23: 39–43. [CrossRef]

Bridges, Robert A., John P. Collins, Erik M. Ferragut, Jason A. Laska, and Blair D. Sullivan. 2015. Multi-level anomaly detection on time-varying graph data. Paper presented at the 2015 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining 2015, Paris, France, August 25–28, pp. 579–83. [CrossRef]

Bun, Joel, Jean-Philippe Bouchaud, and Marc Potters. 2017. Cleaning large correlation matrices: Tools from random matrix theory. *Physics Reports* 666: 1–109. [CrossRef]

Calvo, Clara, Carlos Ivorra, and Vicente Liern. 2014. Fuzzy portfolio selection with non-financial goals: Exploring the efficient frontier. *Annals of Operations Research* 245: 31–46. [CrossRef]

Campbell, Rachel, Ronald Huisman, and Kees Koedijk. 2001. Optimal portfolio selection in a value-at-risk framework. *Journal of Banking & Finance* 25: 1789–804. [CrossRef]

Cappelli, Carmela, Roy Cerqueti, Pierpaolo D'Urso, and Francesca Di Iorio. 2021. Multiple breaks detection in financial interval-valued time series. *Expert Systems with Applications* 164: 113775. [CrossRef]

Coffey, Greg. 2016. Investment Policy Statement: Elements of a Clearly Defined IPS for Non-Profits. Russell Investments Research, April. Available online: https://russellinvestments.com/-/media/files/us/insights/institutions/non-profit/elements-of-a-clearly-defined-ips-for-non-profits-an-update (accessed on 30 June 2022).

Conci, A., and C. Kubrusly. 2017. Distances between sets—A survey. *Advances in Mathematical Sciences and Applications* 26: 1–18.

del Barrio, Eustasio, Evarist Giné, and Carlos Matrán. 1999. Central limit theorems for the Wasserstein distance between the empirical and the true distributions. *The Annals of Probability* 27: 1009–71. [CrossRef]

Deza, Michel Marie, and Elena Deza. 2013. *Encyclopedia of Distances*. Berlin and Heidelberg: Springer. [CrossRef]

Dose, Christian, and Silvano Cincotti. 2005. Clustering of financial time series with application to index and enhanced index tracking portfolio. *Physica A: Statistical Mechanics and Its Applications* 355: 145–51. [CrossRef]

Drożdż, Stanisław, Jarosław Kwapień, and Paweł Oświęcimka. 2021. Complexity in economic and social systems. *Entropy* 23: 133. [CrossRef]

Drożdż, Stanisław, Jarosław Kwapień, Paweł Oświęcimka, Tomasz Stanisz, and Marcin Wątorek. 2020. Complexity in economic and social systems: Cryptocurrency market at around COVID-19. *Entropy* 22: 1043. [CrossRef]

Drożdż, Stanisław, Ludovico Minati, Paweł Oświęcimka, Marek Stanuszek, and Marcin Wątorek. 2019. Signatures of the crypto-currency market decoupling from the forex. *Future Internet* 11: 154. [CrossRef]

Drożdż, Stanisław, Ludovico Minati, Paweł Oświęcimka, Marek Stanuszek, and Marcin Wątorek. 2020. Competition of noise and collectivity in global cryptocurrency trading: Route to a self-contained market. *Chaos: An Interdisciplinary Journal of Nonlinear Science* 30: 023122. [CrossRef]

Drożdż, Stanislaw, Robert Gębarowski, Ludovico Minati, Pawel Oświęcimka, and Marcin Wątorek. 2018. Bitcoin market route to maturity? Evidence from return fluctuations, temporal correlations and multiscaling effects. *Chaos: An Interdisciplinary Journal of Nonlinear Science* 28: 071101. [CrossRef]

Dubuisson, M.-P., and A. K. Jain. 1994. A modified Hausdorff distance for object matching. Paper presented at 12th International Conference on Pattern Recognition, Jerusalem, Israel, October 9–13, pp. 566–68. [CrossRef]

Duffie, Darrell, and Jun Pan. 1997. An overview of value at risk. *The Journal of Derivatives* 4: 7–49. [CrossRef]

Eisler, Zoltán, and János Kertész. 2006. Scaling theory of temporal correlations and size-dependent fluctuations in the traded value of stocks. *Physical Review E* 73: 046109. [CrossRef]

Eiter, Thomas, and Heikki Mannila. 1997. Distance measures for point sets and their computation. *Acta Informatica* 34: 109–33. [CrossRef]

Fama, Eugene F. 1965. The behavior of stock-market prices. *The Journal of Business* 38: 34–105. [CrossRef]

Fastrich, B., S. Paterlini, and P. Winker. 2014. Constructing optimal sparse portfolios using regularization methods. *Computational Management Science* 12: 417–34. [CrossRef]

Fister, Dušan, Matjaž Perc, and Timotej Jagrič. 2021. Two robust long short-term memory frameworks for trading stocks. *Applied Intelligence* 51: 7177–95. [CrossRef]

Fujita, Osamu. 2013. Metrics based on average distance between sets. *Japan Journal of Industrial and Applied Mathematics* 30: 1–19. [CrossRef]

Gardner, Andrew, Jinko Kanno, Christian A. Duncan, and Rastko Selmic. 2014. Measuring distance between unordered sets of different sizes. Paper presented at Conference on Computer Vision and Pattern Recognition, Columbus, OH, USA, June 23–28, pp. 137–43. [CrossRef]

Gilchrist, Warren. 2000. *Statistical Modelling with Quantile Functions*. Boca Raton: Chapman and Hall/CRC. [CrossRef]

Gopikrishnan, P., M. Meyer, L. A. N. Amaral, and H. E. Stanley. 1998. Inverse cubic law for the distribution of stock price variations. *The European Physical Journal B* 3: 139–40. [CrossRef]

Gustafsson, Fredrik. 2001. *Adaptive Filtering and Change Detection*. New York: John Wiley & Sons, Ltd. [CrossRef]

Hawkins, Douglas M. 1977. Testing a sequence of observations for a shift in location. *Journal of the American Statistical Association* 72: 180–86. [CrossRef]

Hawkins, Douglas M., and K. D. Zamba. 2005. A change-point model for a shift in variance. *Journal of Quality Technology* 37: 21–31. [CrossRef]

Hawkins, Douglas M., Peihua Qiu, and Chang Wook Kang. 2003. The changepoint model for statistical process control. *Journal of Quality Technology* 35: 355–66. [CrossRef]

Imbert, Fred, and Thomas Franck. 2020. Dow plunges 10% amid coronavirus fears for its worst day since the 1987 market crash. *CNBC*, March 12.

Iorio, Carmela, Gianluca Frasso, Antonio D'Ambrosio, and Roberta Siciliano. 2018. A P-spline based clustering approach for portfolio selection. *Expert Systems with Applications* 95: 88–103. [CrossRef]

James, Nick, and Max Menzies. 2021a. A new measure between sets of probability distributions with applications to erratic financial behavior. *Journal of Statistical Mechanics: Theory and Experiment* 2021: 123404. [CrossRef]

James, Nick, and Max Menzies. 2021b. Association between COVID-19 cases and international equity indices. *Physica D: Nonlinear Phenomena* 417: 132809. [CrossRef]

James, Nick, and Max Menzies. 2021c. Efficiency of communities and financial markets during the 2020 pandemic. *Chaos: An Interdisciplinary Journal of Nonlinear Science* 31: 083116. [CrossRef]

James, Nick, and Max Menzies. 2022a. Collective correlations, dynamics, and behavioural inconsistencies of the cryptocurrency market over time. *Nonlinear Dynamics* 107: 4001–17. [CrossRef]

James, Nick, and Max Menzies. 2022b. Dual-domain analysis of gun violence incidents in the United States. *Chaos: An Interdisciplinary Journal of Nonlinear Science* 32: 111101. [CrossRef]

James, Nick, and Max Menzies. 2022c. Estimating a continuously varying offset between multivariate time series with application to COVID-19 in the United States. *The European Physical Journal Special Topics* 231: 3419–3426. [CrossRef] [PubMed]

James, Nick, and Max Menzies. 2022d. Global and regional changes in carbon dioxide emissions: 1970–2019. *Physica A: Statistical Mechanics and Its Applications* 608: 128302. [CrossRef]

James, Nick, and Max Menzies. 2022e. Optimally adaptive Bayesian spectral density estimation for stationary and nonstationary processes. *Statistics and Computing* 32: 45. [CrossRef]

James, Nick, and Max Menzies. 2022f. Spatio-temporal trends in the propagation and capacity of low-carbon hydrogen projects. *International Journal of Hydrogen Energy* 47: 16775–84. [CrossRef]

James, Nick, and Max Menzies. 2023a. Distributional trends in the generation and end-use sector of low-carbon hydrogen plants. *Hydrogen* 4: 174–89. [CrossRef]

James, Nick, and Max Menzies. 2023b. Equivalence relations and L^p distances between time series with application to the Black Summer Australian bushfires. *Physica D: Nonlinear Phenomena* 448: 133693. [CrossRef]

James, Nick, Max Menzies, and Georg A Gottwald. 2022. On financial market correlation structures and diversification benefits across and within equity sectors. *Physica A: Statistical Mechanics and Its Applications* 604: 127682. [CrossRef]

James, Nick, Max Menzies, and Howard Bondell. 2021. Understanding spatial propagation using metric geometry with application to the spread of COVID-19 in the United States. *EPL (Europhysics Letters)* 135: 48004. [CrossRef]

James, Nick, Max Menzies, and Howard Bondell. 2022. In search of peak human athletic potential: A mathematical investigation. *Chaos: An Interdisciplinary Journal of Nonlinear Science* 32: 023110. [CrossRef] [PubMed]

James, Nick, Max Menzies, and Jennifer Chan. 2021. Changes to the extreme and erratic behaviour of cryptocurrencies during COVID-19. *Physica A: Statistical Mechanics and Its Applications* 565: 125581. [CrossRef]

James, Nick, Max Menzies, and Kevin Chin. 2022. Economic state classification and portfolio optimisation with application to stagflationary environments. *Chaos, Solitons & Fractals* 164: 112664. [CrossRef]

James, Nick, Max Menzies, James Chok, Aaron Milner, and Cas Milner. 2023. Geometric persistence and distributional trends in worldwide terrorism. *Chaos, Solitons & Fractals* 169: 113277. [CrossRef]

James, Nick, Max Menzies, Lamiae Azizi, and Jennifer Chan. 2020. Novel semi-metrics for multivariate change point analysis and anomaly detection. *Physica D: Nonlinear Phenomena* 412: 132636. [CrossRef]

Jin, Yan, Rong Qu, and Jason Atkin. 2016. Constrained portfolio optimisation: The state-of-the-art Markowitz models. Paper presented at the 5th the International Conference on Operations Research and Enterprise Systems, Rome, Italy, February 23–25, pp. 388–95. [CrossRef]

Khraibani, Hussein, Bilal Nehme, and Olivier Strauss. 2018. Interval estimation of value-at-risk based on nonparametric models. *Econometrics* 6: 47. [CrossRef]

Kocadağlı, Ozan, and Rıdvan Keskin. 2015. A novel portfolio selection model based on fuzzy goal programming with different importance and priorities. *Expert Systems with Applications* 42: 6898–912. [CrossRef]

Koutra, Danai, Neil Shah, Joshua T. Vogelstein, Brian Gallagher, and Christos Faloutsos. 2016. Delta-Con: Principled massive-graph similarity function with attribution. *ACM Transactions on Knowledge Discovery from Data* 10: 1–43. [CrossRef]

Krause, Jochen, and Marc Paolella. 2014. A fast, accurate method for value-at-risk and expected shortfall. *Econometrics* 2: 98–122. [CrossRef]

Laloux, Laurent, Pierre Cizeau, Jean-Philippe Bouchaud, and Marc Potters. 1999. Noise dressing of financial correlation matrices. *Physical Review Letters* 83: 1467–70. [CrossRef]

Lam, Weng Siew, Weng Hoe Lam, and Saiful Hafizah Jaaman. 2021. Portfolio optimization with a mean-absolute deviation-entropy multi-objective model. *Entropy* 23: 1266. [CrossRef]

Lamoureux, Christopher G., and William D. Lastrapes. 1990. Persistence in variance, structural change, and the GARCH model. *Journal of Business & Economic Statistics* 8: 225–34. [CrossRef]

León, Diego, Arbey Aragón, Javier Sandoval, Germán Hernández, Andrés Arévalo, and Jaime Niño. 2017. Clustering algorithms for risk-adjusted portfolio construction. *Procedia Computer Science* 108: 1334–43. [CrossRef]

Li, Bo, and Ranran Zhang. 2021. A new mean-variance-entropy model for uncertain portfolio optimization with liquidity and diversification. *Chaos, Solitons & Fractals* 146: 110842. [CrossRef]

Li, Jiahan. 2015. Sparse and stable portfolio selection with parameter uncertainty. *Journal of Business & Economic Statistics* 33: 381–92. [CrossRef]

Liagkouras, K., and K. Metaxiotis. 2015. Efficient portfolio construction with the use of multiobjective evolutionary algorithms: Best practices and performance metrics. *International Journal of Information Technology & Decision Making* 14: 535–64. [CrossRef]

Liagkouras, K., and K. Metaxiotis. 2018. Handling the complexities of the multi-constrained portfolio optimization problem with the support of a novel MOEA. *Journal of the Operational Research Society* 69: 1609–27. [CrossRef]

Liu, Yanhui, Parameswaran Gopikrishnan, Pierre Cizeau, Martin Meyer, Chung-Kang Peng, and H. Eugene Stanley. 1999. Statistical properties of the volatility of price fluctuations. *Physical Review E* 60: 1390–400. [CrossRef]

Long, H. Viet, H. Bin Jebreen, I. Dassios, and D. Baleanu. 2020. On the statistical GARCH model for managing the risk by employing a fat-tailed distribution in finance. *Symmetry* 12: 1698. [CrossRef]

Lwin, Khin, Rong Qu, and Graham Kendall. 2014. A learning-guided multi-objective evolutionary algorithm for constrained portfolio optimization. *Applied Soft Computing* 24: 757–72. [CrossRef]

Magdon-Ismail, M., A. Atiya, A. Pratap, and Y. Abu-Mostafa. 2003. The maximum drawdown of the Brownian motion. Paper presented at 2003 IEEE International Conference on Computational Intelligence for Financial Engineering, Hong Kong, China, March 20–23, pp. 243–47. [CrossRef]

Mandelbrot, Benoit. 1963. The variation of certain speculative prices. *The Journal of Business* 36: 394–419. [CrossRef]

Mansour, Nabil, Mohamed Sadok Cherif, and Walid Abdelfattah. 2019. Multi-objective imprecise programming for financial portfolio selection with fuzzy returns. *Expert Systems with Applications* 138: 112810. [CrossRef]

Mantegna, Rosario N., H. Eugene Stanley, and Neil A. Chriss. 2000. An introduction to econophysics: Correlations and complexity in finance. *Physics Today* 53: 70. [CrossRef]

Markowitz, Harry. 1952. Portfolio selection. *The Journal of Finance* 7: 77. [CrossRef]

Meghwani, Suraj S., and Manoj Thakur. 2017. Multi-criteria algorithms for portfolio optimization under practical constraints. *Swarm and Evolutionary Computation* 37: 104–25. [CrossRef]

Milhomem, Danilo Alcantara, and Maria José Pereira Dantas. 2020. Analysis of new approaches used in portfolio optimization: A systematic literature review. *Production* 30: e20190144. [CrossRef]

Moody, J., and M. Saffell. 2001. Learning to trade via direct reinforcement. *IEEE Transactions on Neural Networks* 12: 875–89. [CrossRef] [PubMed]

Moreno, Sebastian, and Jennifer Neville. 2013. Network hypothesis testing using mixed kronecker product graph models. Paper presented at IEEE 13th International Conference on Data Mining, Dallas, TX, USA, December 7–10, pp. 1163–68. [CrossRef]

Müllner, Daniel. 2013. Fastcluster: Fast hierarchical, agglomerative clustering routines forRandPython. *Journal of Statistical Software* 53: 1–18. [CrossRef]

Okorie, David Iheke, and Boqiang Lin. 2020. Stock markets and the COVID-19 fractal contagion effects. *Finance Research Letters* 38: 101640. [CrossRef]

Peel, Leto, and Aaron Clauset. 2015. Detecting change points in the large-scale structure of evolving networks. Paper presented at Twenty-Ninth AAAI Conference on Artificial Intelligence, AAAI'15, Austin, TX, USA, January 25–30, pp. 2914–20.

Pessa, Arthur A. B., Matjaz Perc, and Haroldo V. Ribeiro. 2023. Age and market capitalization drive large price variations of cryptocurrencies. *Scientific Reports* 13. [CrossRef]

Podobnik, Boris, Davor Horvatic, Alexander M. Petersen, and H. Eugene Stanley. 2009. Cross-correlations between volume change and price change. *Proceedings of the National Academy of Sciences of the United States of America* 106: 22079–84. [CrossRef]

Prakash, Arjun, Nick James, Max Menzies, and Gilad Francis. 2021. Structural clustering of volatility regimes for dynamic trading strategies. *Applied Mathematical Finance* 28: 236–74. [CrossRef]

Pun, Chi Seng, and Hoi Ying Wong. 2019. A linear programming model for selection of sparse high-dimensional multiperiod portfolios. *European Journal of Operational Research* 273: 754–71. [CrossRef]

Ranshous, Stephen, Shitian Shen, Danai Koutra, Steve Harenberg, Christos Faloutsos, and Nagiza F. Samatova. 2015. Anomaly detection in dynamic networks: A survey. *Wiley Interdisciplinary Reviews: Computational Statistics* 7: 223–47. [CrossRef]

Rosenfeld, Azriel. 1985. Distances between fuzzy sets. *Pattern Recognition Letters* 3: 229–33. [CrossRef]

Ross, Gordon J. 2014. Sequential change detection in the presence of unknown parameters. *Statistics and Computing* 24: 1017–30. [CrossRef]

Ross, Gordan J. 2015. Parametric and nonparametric sequential change detection in R: The cpm package. *Journal of Statistical Software, Articles* 66: 1–20. [CrossRef]

Ross, Gordon J., and Niall M. Adams. 2012. Two nonparametric control charts for detecting arbitrary distribution changes. *Journal of Quality Technology* 44: 102–16. [CrossRef]

Ross, Gordon J., Dimitris K. Tasoulis, and Niall M. Adams. 2013. Sequential monitoring of a Bernoulli sequence when the pre-change parameter is unknown. *Computational Statistics* 28: 463–79. [CrossRef]

Rudin, Walter. 1976. *Principles of Mathematical Analysis*. New York: McGraw-Hill.

Sagi, Omer, and Lior Rokach. 2018. Ensemble learning: A survey. *WIREs Data Mining and Knowledge Discovery* 8: e1249. [CrossRef]

Salah, Hanen Ben, Jan G. De Gooijer, Ali Gannoun, and Mathieu Ribatet. 2018. Mean–variance and mean–semivariance portfolio selection: A multivariate nonparametric approach. *Financial Markets and Portfolio Management* 32: 419–36. [CrossRef]

Sharpe, William F. 1966. Mutual fund performance. *The Journal of Business* 39: 119–38. [CrossRef]

Shaw, Dong X., Shucheng Liu, and Leonid Kopman. 2008. Lagrangian relaxation procedure for cardinality-constrained portfolio optimization. *Optimization Methods and Software* 23: 411–20. [CrossRef]

Shen, Weiwei, Bin Wang, Jian Pu, and Jun Wang. 2019. The kelly growth optimal portfolio with ensemble learning. *Proceedings of the AAAI Conference on Artificial Intelligence* 33: 1134–41. [CrossRef]

Shonkwiler, R. 1989. An image algorithm for computing the Hausdorff distance efficiently in linear time. *Information Processing Letters* 30: 87–89. [CrossRef]

Sigaki, Higor Y. D., Matjaž Perc, and Haroldo V. Ribeiro. 2019. Clustering patterns in efficiency and the coming-of-age of the cryptocurrency market. *Scientific Reports* 9: 1440. [CrossRef] [PubMed]

Soleimani, Hamed, Hamid Reza Golmakani, and Mohammad Hossein Salimi. 2009. Markowitz-based portfolio selection with minimum transaction lots, cardinality constraints and regarding sector capitalization using genetic algorithm. *Expert Systems with Applications* 36: 5058–63. [CrossRef]

Sortino, Frank A., and Robert van der Meer. 1991. Downside risk. *The Journal of Portfolio Management* 17: 27–31. [CrossRef]

Tanaka, Hideo, Peijun Guo, and I.Burhan Türksen. 2000. Portfolio selection based on fuzzy probabilities and possibility distributions. *Fuzzy Sets and Systems* 111: 387–97. [CrossRef]

Tsay, Ruey S. 2010. *Analysis of Financial Time Series*. Wiley Series in Probability and Statistics. Hoboken: John Wiley & Sons, Inc. [CrossRef]

Ullah, Malik Zaka, Fouad Othman Mallawi, Mir Asma, and Stanford Shateyi. 2022. On the conditional value at risk based on the Laplace distribution with application in GARCH model. *Mathematics* 10: 3018. [CrossRef]

Ünlü, Ramazan, and Petros Xanthopoulos. 2021. A reduced variance unsupervised ensemble learning algorithm based on modern portfolio theory. *Expert Systems with Applications* 180: 115085. [CrossRef]

Valenti, Davide, Giorgio Fazio, and Bernardo Spagnolo. 2018. Stabilizing effect of volatility in financial markets. *Physical Review E* 97: 062307. [CrossRef]

Vercher, Enriqueta, José D. Bermúdez, and José Vicente Segura. 2007. Fuzzy portfolio optimization under downside risk measures. *Fuzzy Sets and Systems* 158: 769–82. [CrossRef]

Wang, Fengzhong, Kazuko Yamasaki, Shlomo Havlin, and H. Eugene Stanley. 2006. Scaling and memory of intraday volatility return intervals in stock markets. *Physical Review E* 73: 026117. [CrossRef] [PubMed]

Wang, Shaojie, Shaobo He, Amin Yousefpour, Hadi Jahanshahi, Robert Repnik, and Matjaž Perc. 2020. Chaos and complexity in a fractional-order financial system with time delays. *Chaos, Solitons & Fractals* 131: 109521. [CrossRef]

Wątorek, Marcin, Stanisław Drożdż, Jarosław Kwapień, Ludovico Minati, Paweł Oświęcimka, and Marek Stanuszek. 2021. Multiscale characteristics of the emerging global cryptocurrency market. *Physics Reports* 901: 1–82. [CrossRef]

Xuan, Xiang, and Kevin Murphy. 2007. Modeling changing dependency structure in multivariate time series. Paper presented at the 24th International Conference on Machine Learning—ICML '07, Corvalis, OR, USA, June 20–24, pp. 1055–62. [CrossRef]

Zhao, Pan, and Qingxian Xiao. 2016. Portfolio selection problem with liquidity constraints under non-extensive statistical mechanics. *Chaos, Solitons & Fractals* 82: 5–10. [CrossRef]