Published in final edited form as:

Quant Finance. 2016 ; 16(7): 1053–1067. doi:10.1080/14697688.2015.1101149.

A semiparametric graphical modelling approach for large-scale equity selection

Han Liu, John Mulvey*, and Tianqi Zhao

Department of Operations Research and Financial Engineering, Princeton University, Princeton, NJ 08544, USA

Abstract

We propose a new stock selection strategy that exploits rebalancing returns and improves portfolio performance. To effectively harvest rebalancing gains, we apply ideas from elliptical-copula graphical modelling and stability inference to select stocks that are as independent as possible. The proposed elliptical-copula graphical model has a latent Gaussian representation; its structure can be effectively inferred using the regularized rank-based estimators. The resulting algorithm is computationally efficient and scales to large data-sets. To show the efficacy of the proposed method, we apply it to conduct equity selection based on a 16-year health care stock data-set and a large 34-year stock data-set. Empirical tests show that the proposed method is superior to alternative strategies including a principal component analysis-based approach and the classical Markowitz strategy based on the traditional buy-and-hold assumption.

Keywords

Equity selection; Machine learning; Graphical model; Rebalancing gains; Semiparametric methods; Elliptical copula; Stability selection; Markowitz strategy

1. Introduction

We propose a new stock selection strategy to improve the performance of investment selection tactics by harvesting ‘rebalancing gains’. This concept has been well documented in both theory and practice. There has been a large literature (Mulvey *et al.* 2001, Mulvey *et al.* 2004, Mulvey *et al.* 2007, Mulvey and Kim 2009, Dempster *et al.* 2008, Dempster *et al.* 2010, Luenberger 2013) that shows the benefits of rebalancing a portfolio as compared with the traditional buy-and-hold rule embedded in the classical Markowitz portfolio model. In particular, we can evaluate the level of rebalancing gains when stocks are approximated as correlated geometric Brownian motions.

The most desirable environment for achieving rebalancing gains requires several components. First, transaction costs (including market impact costs) must be relatively small, for example, via actively traded futures markets. Second and importantly, the returns

*Corresponding author. mulvey@princeton.edu.

Disclosure statement

No potential conflict of interest was reported by the authors.

of the instruments should be as independent as possible. Third, the instruments must display a positive drift. As an idealized example, suppose that we are able to discover 10 investment tactics with the same growth level ρ , covariance matrix Σ , and independent returns (as specified by white noise terms). Ignoring the transaction costs, rebalancing gains are strictly a function of the volatility of the individual instruments. Table 1 shows the level of rebalancing gains for several levels of volatility (Luenberger 2013).

Clearly, rebalancing gains are ideally suited to markets possessing low correlation, relatively independent returns and high volatility. To some degree, these characteristics are part and parcel of commodity markets (Gorton and Rouwenhorst 2004, Higgs and Worthington 2008, Gorton *et al.* 2013, Mulvey 2012).

There are several issues regarding practical implementation of a rebalancing strategy (Mulvey and Kim 2009). The first issue is to find a set of relatively independent assets. This task is challenging and unlikely to be perfectly achieved in the real world since financial instruments are correlated. For example, all the stocks in the market are exposed to a degree of systemic risk, and exhibit co-movements with the overall market. As a second issue, correlations change over time. For example, financial instruments exhibit higher levels of correlation during downside regimes than upside regimes.

To handle these challenges, we propose a new machine learning framework to construct equity portfolios whose members are as independent as possible. Our method exploits new ideas of elliptical-copula modelling, latent Markovian graphical models and stability selection.

Copula models are popular in quantitative finance (Embrechts *et al.* 2003, Cherubini *et al.* 2004, Costa Dias 2004, Dias and Embrechts 2004). They are particularly useful for modelling dependency structures among financial instruments. For example, the Gaussian copula has attracted much interest because of its simplicity and mathematical elegance (Li 2000). However, its flaw in not being able to accurately model tail dependence has been discovered during financial crises. In this paper, we exploit the more realistic elliptical-copula distribution family to model stock data. Unlike the Gaussian copula model, the elliptical-copula model is more general and can depict nontrivial tail dependence. In the statistical machine learning literature, the elliptical-copula model has been shown to be equivalent to the ‘transelliptical’ model (Details are provided in section 2.3).

We define a notion of an elliptical-copula graphical model which provides a graphical representation of multivariate elliptical distribution. To interpret the graph, we develop a hierarchical Gaussian representation for any elliptical distribution. Such a latent Gaussian representation allows us to interpret the elliptical graphical model as a latent Markov random field. It also provides a formal mechanism to select a subset of stocks that are ‘as independent as possible’.

We exploit a regularized rank-based estimator to depict the structure of the latent Gaussian Markov random field of a large number of stock returns. This graph encodes the conditional independence relationship of the latent Gaussian random variables. Next, we propose an efficient path-following algorithm that obtains a sequence of graph estimates corresponding

to different sparsity levels. This algorithm is then embedded within a stability selection framework to obtain a subset of stocks. The selected stocks form an equally weighted portfolio for investment.

We apply the proposed method on two stock data-sets, showing that it significantly outperforms the classical buy-and-hold rule that underlies the Markowitz strategy. In addition, our graphical model outperforms a principal component analysis-based strategy.

The rest of the paper is organized as follows: section 2 describes the proposed stock selection method. In particular, section 2.3 introduces the elliptical-copula graphical model and its hierarchical interpretation. Sections 2.4 and 2.5 describe the estimator and graph estimation algorithm. Section 2.6 introduces the stability selection method, and section 2.7 describes the complete algorithm. In section 3, we provide some theoretical guarantees of our stock selection procedure. In section 4, we apply the proposed method to stock data and evaluate the performance of a rebalancing strategy compared with other strategies, showing the potential advantages of the proposed selection procedure.

Notation

Capital italic letters denote random variables (e.g. return of one stock X_1) and bold italic letters to denote random vectors (e.g. a random vector is denoted as $\mathbf{X}$). All matrices are denoted by normal bold letters (e.g. $\mathbf{A}$). For any square matrix $\mathbf{A}$ and $\mathbf{B}$, $\|\mathbf{A}\|_1 = \sum_{j,k} |\mathbf{A}_{jk}|$ is

the elementwise ℓ_1 -norm, $\|\mathbf{A}\|_F = \sqrt{\sum_{j,k} \mathbf{A}_{jk}^2}$ is the Frobenius norm and $\langle \mathbf{A}, \mathbf{B} \rangle = \text{Tr}(\mathbf{A}^T \mathbf{B})$ is the matrix inner product. Let $\mathbb{I}\{\cdot\}$ be the indicator function.

2. A graphical model approach for selection

We employ new tools from probabilistic graphical models and stability selection to construct a portfolio that can effectively exploit rebalancing returns.

More specifically, we adopt the elliptical-copula graphical models for multivariate stock selection and develop efficient computational algorithms for parameter learning and graph estimation.

2.1. Preliminaries on undirected graphical models

The undirected graphical model is a powerful tool for modelling the relationships among a large number of random variables. In this subsection, we introduce a simplest notion of undirected graphical model—Gaussian graphical model. Next, we extend the approach to the more sophisticated elliptical-copula graphical model.

Let $\mathbf{X} = (X_1, \dots, X_d)^T \in \mathbb{R}^d$ be a d -dimensional Gaussian random vector $\mathbf{X} \sim \mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$. Define an undirected graph $G = (V, E)$ with the vertex set $V := \{X_1, \dots, X_d\}$ and edge encoded in the adjacency matrix $E \in \{0, 1\}^{d \times d}$. We denote $E_{jk} = 1$ if there is an edge connecting X_j and X_k , and $E_{jk} = 0$ otherwise. $\mathbf{X}$ is Markov to G if $E_{jk} = 0$ if and only if the variables X_j and X_k are conditionally independent given the rest of variables, denoted by a subvector $\mathbf{X}_{\setminus\{j,k\}}$. Dempster (1972) shows that under the Gaussian distribution, the graph G is completely

characterized by the sparsity pattern of the inverse covariance matrix, also known as the precision matrix $\Theta = \Sigma^{-1}$. In particular, $E_{jk} = 0$ if and only if $\Theta_{jk} = 0$.

The Gaussian graphical model, though simple and elegant, is unrealistic for our purposes. Specifically, stock returns are known to be heavy-tailed with nontrivial tail dependence, while the Gaussian distribution fails these properties. For example, in downside regimes, stock returns tend to drop drastically as many investors simultaneously seek refuge in safer assets such as cash or bonds, whereas in upside regimes, such collective behaviour is less likely to be observed and stock returns are more stable. In the next subsection, we relax the Gaussian modelling assumptions.

2.2. Elliptical copula distribution family

A *copula* is a joint cumulative distribution function of random variables that have uniform marginal distributions. For any random vector $\mathbf{X} \in (X_1, \dots, X_d)^T \in \mathbb{R}^d$ with marginal distribution functions $F_1, \dots, F_d$, a key observation is that the transformed random variable $F_j(X_j)$ has a uniform distribution on $[0, 1]$ for any $j = 1, \dots, d$. We define the copula for $\mathbf{X}$ as a function $C: [0, 1]^d \mapsto [0, 1]$ such that

$$C(u_1, u_2, \dots, u_d) := \mathbb{P}(F_1(X_1) \leq u_1, \dots, F_d(X_d) \leq u_d).$$

Sklar's theorem shows that the joint distribution function F of $\mathbf{X}$ can be expressed using its copula function and marginal distributions, i.e.

$$F_{X_1, \dots, X_d}(x_1, \dots, x_d) = C(F_1(x_1), \dots, F_d(x_d)). \quad (2.1)$$

Equation (2.1) shows that any joint multivariate distribution contains two ingredients: the marginal distribution and the copula function. In particular, the marginal distributions capture the marginal information of each individual component in $\mathbf{X}$, while the copula function captures the interrelationships among these components.

The concept of copula is not new. It has been extensively employed for market data calibration, stress tests, credit risk evaluation, derivatives pricing, etc. However, the way we apply copula to portfolio analysis is new: we apply the elliptical-copula distribution family to model the stock returns and define graphical models based on the elliptical copula.

To rigorously define elliptical-copula distribution family, we start with a formal definition of an elliptical distribution. To this end, we define $\mathbf{Y}$ to be elliptically distributed if its characteristic function $\Psi_{\mathbf{Y}}$ can be written as

$$\Psi_{\mathbf{Y}}(\mathbf{t}) := \mathbb{E}[e^{it^T \mathbf{X}}] = \exp(it^T \boldsymbol{\mu}) \psi(\mathbf{t}^T \boldsymbol{\Sigma} \mathbf{t}),$$

where $\boldsymbol{\mu}$ is the location parameter (also the mean if it exists), $\boldsymbol{\Sigma}$ is a positive definite scatter matrix, and $\psi(\cdot)$ is specific to each elliptical random variable and is often called the

‘characteristic generator’. Hence an elliptical distribution is uniquely determined by the above three ingredients, and we denote

$$\mathbf{Y} \sim E_d(\boldsymbol{\mu}, \boldsymbol{\Sigma}, \psi),$$

for any random vector $\mathbf{Y}$ that follows an elliptical distribution. The density of $\mathbf{Y}$ has the following form: $f_{\mathbf{Y}}(\mathbf{y}) \propto h(\mathbf{y} - \boldsymbol{\mu})^T \boldsymbol{\Sigma}^{-1}(\mathbf{y} - \boldsymbol{\mu})$, where h is uniquely determined by ψ .

Significantly, the elliptical distribution family strictly extends the Gaussian family and contains many heavy-tailed distributions, for example, multivariate t-distribution, Cauchy, logistic, Kotz, symmetric Pearson type-II and type-VII distributions. Intuitively, the elliptical distribution family includes all distributions whose density function has elliptical contours.

Next, we introduce the elliptical-copula distribution family:

Definition 2.1 (Elliptical Copula Distribution Family)—The elliptical-copula distribution family contains all distributions with elliptical-copulas.

Any distribution $\mathbf{X}$ in the elliptical-copula family is generated by combining an elliptical-copula and possibly non-elliptical marginal distributions. In fact, Liu *et al.* (2012b) proved that for any $\mathbf{X} = (X_1, \dots, X_d)$ that belongs to the elliptical-copula family, there is a set of latent variables $\mathbf{Y} = (Y_1, \dots, Y_d)$ corresponding to $\mathbf{X}$ with an elliptical distribution and zero mean. Concretely, there exists a set of strictly increasing functions $g_1, \dots, g_d$ such that $(g_1(X_1), \dots, g_d(X_d)) \sim E_d(\mathbf{0}, \boldsymbol{\Sigma}_0, \psi)$ where $\text{diag}(\boldsymbol{\Sigma}_0) = \mathbf{I}_d$ so that the model is identifiable. Hence the elliptical-copula family is also called the transelliptical family, and can be represented as

$$\mathbf{X} \sim TE_d(\boldsymbol{\Sigma}_0, \psi; g_1, \dots, g_d).$$

Here, $\boldsymbol{\Sigma}_0$ is the latent correlation matrix, as it is the correlation matrix of the latent random vector $\mathbf{Y}$.

Modelling stock returns via elliptical copulae has many advantages: the elliptical-copula distribution can naturally depict heavy-tailed, asymmetric data and tail dependencies—an important property of stock return data.

2.3. Elliptical copula graphical model and its latent Gaussian representation

Let $\mathbf{X} = (X_1, \dots, X_d)^T$ be an elliptical-copula distribution, denoted by $\mathbf{X} \sim TE_d(\boldsymbol{\Sigma}_0, \psi; g_1, \dots, g_d)$. We define a graph G based on the sparsity pattern of the latent inverse correlation matrix $\boldsymbol{\Theta} := \boldsymbol{\Sigma}_0^{-1}$. There is an edge connecting nodes X_j and X_k if and only if $\boldsymbol{\Theta}_{jk} \neq 0$. Given any semiparametric transelliptical distribution, such a graph G is always well defined. The key question is: how shall we interpret the obtained graph? Herein, we present a hierarchical latent Gaussian representation of the elliptical-copula distribution. Such a representation allows us to leverage the existing results developed under the Gaussian graphical model.

The elliptical-copula graphical model can be viewed as a three-layer generative model (figure 1). The first layer encompasses all the observed elliptical-copula (or transelliptical) random variables. The second layer includes a set of latent elliptical variables, each of which can generate a corresponding observed transelliptical variable by marginal monotone transformation. The third layer includes a set of Gaussian random variables which could generate the second layer elliptical random variables by global stochastic scaling. In particular, the three layers share the same scatter matrix Σ_0 , thus they have the same graph structure (since the graph is completely endowed by the sparsity pattern of Σ_0^{-1}). In the following, we provide more detailed interpretations of the graph at different layers.

- We start from the second layer—the latent elliptical variables. By definition, G is the sparsity pattern of the latent precision matrix Θ , which is the inverse of the correlation matrix Σ_0 of the latent random vector $\mathbf{Y}$, where $Y_j = g_j^{-1}(X_j)$ for $j = 1, \dots, d$. Hence G encodes the correlation information of the latent variables in $\mathbf{Y}$. By definition, $\mathbf{Y} \sim E_d(\mathbf{0}, \Sigma_0, \psi)$. In the proof of Lemma 3.7 in Liu *et al.* (2013), it is shown that

$$\mathbf{Y}_J | \mathbf{Y}_{J^c} \sim E_{|J|}(\tilde{\boldsymbol{\mu}}, [\Theta_{JJ}]^{-1}, \tilde{\psi}). \quad (2.2)$$

for some $\tilde{\boldsymbol{\mu}} \in \mathbb{R}^{|J|}$ and $\tilde{\psi}$. As Θ_{JJ} is diagonal if and only if $[\Theta_{JJ}]^{-1}$ is diagonal, thus Θ_{JJ} being diagonal is equivalent to $\mathbf{Y}_J$ being pairwise uncorrelated given $\mathbf{Y}_{J^c}$. When $J = \{j, k\}$, this implies that $\Theta_{jk} = 0$ if and only if Y_j and Y_k are uncorrelated conditional on $\mathbf{Y}_{\setminus\{j,k\}}$. Hence we have $E_{jk} = \mathbb{I}\{\text{Corr}(Y_j, Y_k | \mathbf{Y}_{\setminus\{j,k\}}) = 0\}$.

- We next consider the first layer—the observed transelliptical variables. For the observed random variables, Σ_0 does not directly characterize the correlation between random variables in $\mathbf{X}$. This is because correlation matrix is not invariant under marginal transformation. Instead, we consider the following rank-based correlation statistic:

$$\tau_{jk} := \mathbb{P}\left((X_j - \tilde{X}_j)(X_k - \tilde{X}_k) > 0\right) - \mathbb{P}\left((X_j - \tilde{X}_j)(X_k - \tilde{X}_k) < 0\right) \quad (2.3)$$

where $\tilde{\mathbf{X}} = (\tilde{X}_1, \dots, \tilde{X}_d)$ is an independent copy of $\mathbf{X}$. τ_{jk} is the population version of Kendall's tau statistic. It is not difficult to see that τ_{jk} is invariant to monotone marginal transformation, i.e. $\tau_{jk} = \mathbb{P}((Y_j - \tilde{Y}_j)(Y_k - \tilde{Y}_k) > 0) - \mathbb{P}((Y_j - \tilde{Y}_j)(Y_k - \tilde{Y}_k) < 0)$, which shows that the rank correlation of the observed random variable is the same as that of the latent elliptical random variables. Linsdskog *et al.* (2003) and Liu *et al.* (2013) showed that τ_{jk} has the following simple relationship with $[\Sigma_0]_{jk}$ for elliptically distributed random variables: $\tau_{jk} = (2/\pi) \arcsin([\Sigma_0]_{jk})$. Such a relationship also holds for the conditional version:

$$\tau(X_j, X_k | \mathbf{X}_{\setminus\{j,k\}}) = \tau(Y_j, Y_k | \mathbf{Y}_{\setminus\{j,k\}}) = \frac{2}{\pi} \arcsin \left(\left[[\Theta_{\{j,k\}}]^{-1} \right]_{jk} \right),$$

where we employed the fact that $\mathbf{Y}_{\{j,k\}} | \mathbf{Y}_{\setminus\{j,k\}} \sim E_2(\tilde{\boldsymbol{\mu}}, [\Theta_{\{j,k\}}]^{-1}, \tilde{\boldsymbol{\psi}})$ given by (2.2). Thus, $\Theta_{jk} = 0$ if and only if X_j and X_k are conditionally uncorrelated under rank association given $\mathbf{X}_{\setminus\{j,k\}}$, i.e. $E_{jk} = \mathbb{I}\{\tau(X_j, X_k | \mathbf{X}_{\setminus\{j,k\}}) = 0\}$.

- Lastly, we consider the third layer—the latent Gaussian variables. For any elliptically distributed random vector $\mathbf{Y} \sim E_d(\boldsymbol{\mu}, \boldsymbol{\Sigma}, \boldsymbol{\psi})$, it has a latent Gaussian representation (Liu *et al.* 2012b). Concretely, suppose $\boldsymbol{\Sigma} = \mathbf{A}\mathbf{A}^T$, then we have

$$\mathbf{Y} \stackrel{d}{=} \boldsymbol{\mu} + \eta \mathbf{Z} / \|\mathbf{A}^* \mathbf{Z}\|_2,$$

where $\mathbf{Z} \sim N_d(\mathbf{0}, \mathbf{I})$. $\eta = 0$ is a random variable that is independent of $\mathbf{Z}$, and $\mathbf{A}^*$ is the Moore-Penrose pseudoinverse of $\mathbf{A}$. Hence the third level is composed of the latent random variables $\mathbf{Z} = (Z_1, \dots, Z_d)$. By conditional Gaussian formula, we have that $\mathbf{Z}_J | \mathbf{Z}_{J^c} \sim N(\boldsymbol{\Sigma}_{JJ^c} [\boldsymbol{\Sigma}_{J^c J^c}]^{-1} \mathbf{Z}_{J^c}, [\boldsymbol{\Sigma}_{JJ}]^{-1})$, therefore $\Theta_{jk} = 0$ if and only if Z_j and Z_k are conditionally uncorrelated given $\mathbf{Z}_{\setminus\{j,k\}}$, and, by the Gaussian property, if and only if they are conditionally independent. Hence $E_{jk} = \mathbb{I}\{Z_j \perp Z_k | \mathbf{Z}_{\setminus\{j,k\}}\}$.

This hierarchical framework characterizes the interdependency structure between stocks. Let $\mathbf{X}_1, \dots, \mathbf{X}_n$ denote the observed shistorical returns of d stocks across n periods. To construct an ‘optimal’ rebalancing portfolio, we seek to select stocks that are ‘as independent as possible’. For this, we first estimate the undirected elliptical-copulal graph that underlies d random variables (stock returns), and then select a set of stocks each of which is ‘isolated’ from the rest of the graph. The estimation and selection procedure is described in the next subsection. Notably, our method scales with very large dimension d . This feature is crucial, as it allows us to choose from a large pool of stocks, so the selected stocks could be very different from each other and potentially close to being independent.

2.4. Rank-based estimator

As G is defined by the sparsity pattern of Θ , our graph estimating procedure is based on three steps: we first estimate the latent correlation matrix using an estimator $\hat{\boldsymbol{\Sigma}}$, then solve for a sparse approximation $\hat{\Theta}$ of the inverse of $\hat{\boldsymbol{\Sigma}}$, and lastly we estimate G via $\hat{\Theta}$ and a stability-based method. In this and the following two sections, we describe the three steps, respectively.

From the previous subsection, it is not difficult to see that the rank-based Kendall’s tau statistic is a good estimator for $\boldsymbol{\Sigma}_0$, because of its two desirable properties: (i) it is invariant under the monotone marginal transformation, so that the observed $\mathbf{X}$ and the latent variables in $\mathbf{Y}$ share the same Kendall’s tau statistic; (ii) it has a simple relationship with the

correlation matrix under the elliptical distribution, i.e. $[\boldsymbol{\Sigma}_0]_{jk} = \sin(\frac{\pi}{2} \tau_{jk})$. Liu *et al.* (2012a) and Xue and Zou (2012) exploited the same estimator for inferring the latent correlation

matrix Σ in the Gaussian copula family (also known as nonparanormal family), which is an extension of the Gaussian distribution using marginal monotone transformation. It was further extended to the elliptical-copula (transelliptical) family by Han and Liu (2012) and Liu *et al.* (2012b).

Formally, our estimator is defined as follows:

$$\hat{\Sigma}_{jk} = \sin \left(\frac{\pi}{2} \hat{\tau}_{jk} \right), \quad (2.4)$$

where

$$\hat{\tau}_{jk} := \frac{2}{n(n-1)} \sum_{1 \leq i < i' \leq n} [\mathbb{I}\{(X_{ij} - X_{i'j})(X_{ik} - X_{i'k}) > 0\} - \mathbb{I}\{(X_{ij} - X_{i'j})(X_{ik} - X_{i'k}) < 0\}] \quad (2.5)$$

is the estimated Kendall's tau statistic in the replacement of the population version (2.3).

Liu *et al.* (2012a) showed that $\hat{\Sigma}$ is a consistent estimator of the latent correlation matrix Σ (with respect to elementwise sup norm $\|\cdot\|_{\max}$), even when the order of d is nearly exponentially larger than n . Concretely, we have with probability at least $1 - 1/\max\{d, n\}$,

$$\|\hat{\Sigma} - \Sigma_0\|_{\max} \leq 2.45\pi \sqrt{\frac{\log d}{n}}.$$

We next introduce a statistical procedure that efficiently estimate the precision matrix Θ by plugging in the Kendall's tau estimate $\hat{\Sigma}$. With a properly chosen tuning parameter λ , the estimated graph has a sparse representation, which facilitates the stock selection procedure. Moreover, when the true graph is indeed sparse, the estimated graph is also consistent. Thus, our selection procedure has a solid statistical foundation.

2.5. Smooth-projected neighbourhood pursuit algorithm for graph estimation

We next introduce a statistical procedure that efficiently estimates sparsity of the precision matrix Θ using the Kendall's tau estimate $\hat{\Sigma}$. Note that although Θ is the inverse of Σ , naively inverting $\hat{\Sigma}$ is not a good procedure, as $\hat{\Sigma}$ is not invertible when dimension d is larger than n , and even if it is, its inverse is not sparse.

There are many procedures for transelliptical graph estimation using the Kendall's statistics. We adopt the neighbourhood pursuit method introduced by Meinshausen and Bühlmann (2006). The intuition originates from the Gaussian graphical model. Recall from §2.3 that for Gaussian random variables, $Z_j | Z_{\setminus j} \sim \mathcal{N}(\Sigma_{j,\setminus j} [\Sigma_{\setminus j,\setminus j}]^{-1} Z_{\setminus j}, [\Theta_{jj}]^{-1})$. Therefore, we have

$$\mathbb{E}[Z_j | \mathbf{Z}_{\setminus j}] = \Sigma_{j, \setminus j} [\Sigma_{\setminus j, \setminus j}]^{-1} \mathbf{Z}_{\setminus j} = \mathbb{E}[\boldsymbol{\beta}^T \mathbf{Z}_{\setminus j}], \quad (2.6)$$

where $\boldsymbol{\beta} = [\Sigma_{\setminus j, \setminus j}]^{-1} \Sigma_{\setminus j, j}$ (2.6) is exactly a regression problem. Moreover, by matrix block inverse formula, we have $\boldsymbol{\beta} = \boldsymbol{\Theta}_{\setminus j, j} / \boldsymbol{\Theta}_{jj}$. This shows that $\boldsymbol{\beta}$ encodes the sparsity pattern of $\boldsymbol{\Theta}_{\setminus j, j}$. Hence the neighbourhood pursuit algorithm is based on regressing Z_j onto other variables using an ℓ_1 penalized minimization method known as LASSO regression: let $\hat{\mathbf{R}} \in \mathbb{R}^{d \times d}$, then we solve

$$\hat{\mathbf{R}}_{\setminus j, j} = \underset{\boldsymbol{\beta} \in \mathbb{R}^{d-1}}{\operatorname{argmin}} \frac{1}{n} \|\mathbf{Z}_{\cdot, j} - \mathbf{Z}_{\cdot, \setminus j} \boldsymbol{\beta}\|_2^2 + \lambda \|\boldsymbol{\beta}\|_1,$$

where $\mathbf{Z} = (\mathbf{Z}_1, \dots, \mathbf{Z}_n)^T \in \mathbb{R}^{n \times d}$ is the design matrix and λ is a tuning parameter. The ℓ_1 -regularization is imposed so that the solution has a sparse representation, a common variable selection method used in statistics. Solving the above problems for all j and letting $\operatorname{diag}(\hat{\mathbf{R}}) = \mathbf{0}$, we derive an estimator $\hat{\mathbf{R}}$ for $\boldsymbol{\Theta}$.

If we expand the objective function, we find that it can be rewritten as follows:

$$\hat{\mathbf{R}}_{\setminus j, j} = \underset{\boldsymbol{\beta} \in \mathbb{R}^{d-1}}{\operatorname{argmin}} \boldsymbol{\beta}^T \hat{\mathbf{S}}_{\setminus j, \setminus j} \boldsymbol{\beta} - 2 \hat{\mathbf{S}}_{\setminus j, j} \boldsymbol{\beta} + \lambda \|\boldsymbol{\beta}\|_1.$$

where $\hat{\mathbf{S}} = \mathbf{Z}^T \mathbf{Z} / n$ is the Pearson sample covariance matrix. To solve for the elliptical-copula (transelliptical) graph, a natural idea is to replace $\hat{\mathbf{S}}$ by the Kendall's tau covariance matrix $\hat{\boldsymbol{\Sigma}}$. Although statistically this method enjoys good property, it is computationally very challenging: $\hat{\boldsymbol{\Sigma}}$ is possible to be not positive definite, and when it is not, the above minimization with $\hat{\boldsymbol{\Sigma}}$ plugged in is not a convex problem, which is very difficult to solve.

To handle this challenge, we apply the smooth-projected neighbourhood pursuit algorithm proposed by Zhao, Roeder *et al.* (2012) for estimating Gaussian copula (nonparanormal) graphical models. In this algorithm, we first obtain a positive definite estimator $\tilde{\boldsymbol{\Sigma}}$ which is the projection of the Kendall's tau estimate $\hat{\boldsymbol{\Sigma}}$. Ideally, we want to project the matrix under the elementwise max norm $\|\cdot\|_{\max}$. However, the projection problem under this norm is difficult to compute due to the non-smooth nature of the max norm. Therefore, we first define a smooth norm that approximates the elementwise max norm: for any matrix $\mathbf{A}$, define

$$\|\mathbf{A}\|_{\infty}^{\rho} = \max_{\|\mathbf{B}\|_1 \leq 1} \left\{ \langle \mathbf{B}, \mathbf{A} \rangle - \frac{\rho}{2} \|\mathbf{B}\|_F^2 \right\},$$

In particular, when $\|\cdot\|_{\infty}^{\rho}$ simply becomes $\|\cdot\|_{\max}$ by norm duality. We then define the smooth projection $\tilde{\boldsymbol{\Sigma}}$ of $\hat{\boldsymbol{\Sigma}}$ as

$$\tilde{\Sigma} = \underset{\Sigma \in S\mathbb{R}_+^d}{\operatorname{argmin}} \|\hat{\Sigma} - \Sigma\|_\infty^\rho, \quad (2.7)$$

where $S\mathbb{R}_+^d$ denotes the set of symmetric positive definite $d \times d$ matrices. Now that $\tilde{\Sigma}$ is positive definite and is an approximation of $\hat{\Sigma}$, we can define the estimate $\hat{\mathbf{R}}$ by solving

$$\hat{\mathbf{R}}_{\setminus j,j} = \underset{\beta \in \mathbb{R}^{d-1}}{\operatorname{argmin}} \beta^T \tilde{\Sigma}_{\setminus j,\setminus j} \beta - 2\tilde{\Sigma}_{\setminus j,j} \beta + \lambda \|\beta\|_1, \quad (2.8)$$

for all $j = 1, \dots, d$ and setting $\operatorname{diag}(\hat{\mathbf{R}}) = \mathbf{0}$.

The rest of this subsection is devoted to finding efficient algorithms for solving (2.7) and (2.8). The optimization problem (2.7) can be solved by the fast proximal gradient algorithm (Zhao, Roeder *et al.* 2012). We present the procedure in algorithm 1.

Algorithm 1

Fast proximal gradient descent algorithm for smooth projection

-
- 1: $\hat{\Sigma} \leftarrow \text{Proximal-Gradient}(\hat{\Sigma}, \rho)$
 - 2: **input:** A symmetric matrix $\hat{\Sigma}$, a smoothing parameter ρ
 - 3: **output:** A positive definite matrix $\hat{\Sigma}$, as the solution of (2.7).
 - 4: **initialize:** $t \leftarrow 1$, $\mathbf{U}^{(0)} = \mathbf{V}^{(0)} = \hat{\Sigma}^{(0)}$
 - 5: **repeat**
 - 6: $\gamma_t \leftarrow 2/(1+t)$
 - 7: $\mathbf{U}^{(t)} \leftarrow (1 - \gamma_t)\hat{\Sigma}^{(t-1)} + \gamma_t \mathbf{V}^{(t-1)}$
 - 8: $\mathbf{G}^{(t)} \leftarrow \nabla \|\hat{\Sigma} - \mathbf{U}\|_\infty^\rho|_{\mathbf{U}=\mathbf{U}^{(t)}}$, where the gradient can be calculated by (2.9)
 - 9: $\mathbf{V}^{(t)} \leftarrow \Pi_+ \left(\mathbf{V}^{(t-1)} - \frac{\rho}{\gamma_t} \mathbf{G}^{(t)} \right)$, where $\Pi_+(\cdot)$ can be calculated by (2.10)
 - 10: $\hat{\Sigma}^{(t)} \leftarrow (1 - \gamma_t)\hat{\Sigma}^{(t-1)} + \gamma_t \mathbf{V}^{(t)}$
 - 11: **until** Converge
 - 12: **return** $\hat{\Sigma} \leftarrow \hat{\Sigma}^{(t)}$
-

In particular, Step 8 in algorithm 1 requires to calculate the gradient of the smoothed elementwise ℓ_∞ -norm. An explicit calculation gives that

$$\nabla \|\hat{\Sigma} - \mathbf{U}\|_\infty^\rho = -\tilde{\mathbf{T}}, \quad (2.9)$$

where

$$\tilde{\mathbf{T}}_{jk} = \text{sign}(\hat{\Sigma}_{jk} - \mathbf{U}_{jk}) \cdot \max \left\{ \left| \frac{\hat{\Sigma}_{jk} - \mathbf{U}_{jk}}{\rho} \right| - \theta, 0 \right\}$$

and θ is the minimum non-negative constant such that $\|\tilde{\mathbf{U}}\|_1 = 1$. $\mathbf{T}$ is essentially a ‘soft thresholding’ function, which means that it thresholds the value of $\hat{\Sigma}_{jk} - \mathbf{U}_{jk}$ as 0 if its absolute value is less than $\rho \cdot \theta$. A pictorial explanation of the soft thresholding function is provided in figure 2.

Moreover, Step 9 in algorithm 1 needs to solve the following positive definite projection problem under the Frobenius norm:

$$\Pi_+(\mathbf{A}) = \underset{\mathbf{B} \in \mathbb{S}_{+}^{\mathbb{R}^d}}{\text{argmin}} \|\mathbf{B} - \mathbf{A}\|_F^2,$$

for any symmetric matrix $\mathbf{A} \in \mathbb{R}^{d \times d}$. Lemma 3.2 in Zhao, Roeder *et al.* (2012) showed that $\Pi_+(\mathbf{A})$ has a closed-form solution: As $\mathbf{A}$ is symmetric, it has the eigenvalue decomposition as $\mathbf{A} = \sum_{j=1}^d \sigma_j \mathbf{v}_j \mathbf{v}_j^T$, where σ_j ’s are the eigenvalues and $\mathbf{v}_j$ are the eigenvectors. The solution to the projection problem under Frobenius norm is

$$\Pi_+(\mathbf{A}) = \sum_{j=1}^d \max(\sigma_j, 0) \mathbf{v}_j \mathbf{v}_j^T. \quad (2.10)$$

This shows that the projection to the space of positive definite matrix under Frobenius norm is simply a thresholding of the eigenvalues.

The following theorem guarantees fast convergence of algorithm 1:

Theorem 2.2 (Zhao, Roeder *et al.* 2012)—*To achieve the accuracy ε such that $\|\hat{\Sigma} - \Sigma^{(t)}\|_\infty^\rho - \|\hat{\Sigma} - \tilde{\Sigma}\|_\infty^\rho \leq \varepsilon$, we need at most the following number of iterations:*

$$t = \sqrt{2\|\Sigma^{(0)} - \tilde{\Sigma}\|_F^2 / (\rho\varepsilon)} - 1 = O(\sqrt{1/(\rho\varepsilon)}).$$

This convergence rate is actually optimal in the sense that it is unimprovable over the class of all gradient-based computational algorithms (Nesterov and Nesterov 2004).

The minimization problem (2.8) can be efficiently solved using existing solvers. For example, it has the special structure of a ℓ_1 -penalized regression (LASSO) and thus can be solved using many optimization algorithms, such as the coordinate descent (Friedman *et al.* 2007). Alternatively, we can use general convex programming algorithms such as the interior point method (Vanderbei 1999).

Once $\hat{\mathbf{R}}$ is obtained, we calculate the graph estimator $\hat{G} = (V, \hat{E})$, in which a pair (j, k) satisfies that $\hat{E}_{jk} = 0$ if and only if $\hat{\mathbf{R}}_{jk} = 0$. Algorithm 2 summarizes the approach. For each λ , this algorithm outputs the estimated graph that underlies the stock returns. Choosing a large λ will impose large penalization of the ℓ_1 -norm of the $\hat{\Theta}$ and result in a sparse graph and vice versa.

Algorithm 2

Graph Estimation

-
- 1: $\hat{G}^\lambda \leftarrow \text{Graph-Estimation}(X_1, \dots, X_n; \lambda)$
 - 2: **input:** Historical return of d stocks over n periods $X_1, \dots, X_n \in \mathbb{R}^d$, tuning parameter λ
 - 3: **output:** Estimated graph $\hat{G}^\lambda$ corresponding to the tuning parameter λ
 - 4: Calculate Kendall's tau correlation matrix by (2.4) and (2.5).
 - 5: Let $\tilde{\Sigma} \leftarrow \text{Proximal-Gradient}(\tilde{\Sigma})$ as in algorithm 1
 - 6: Obtain $\hat{\mathbf{R}}$ by solving

$$\hat{\mathbf{R}} = \underset{\mathbf{R} \in \mathbb{R}^{d \times d}, \text{diag}(\mathbf{R}) = \mathbf{0}}{\text{argmin}} \sum_{j=1}^d \left\{ \mathbf{R}_{\setminus j, j}^T \tilde{\Sigma}_{\setminus j, \setminus j} \mathbf{R}_{\setminus j, j} - 2 \tilde{\Sigma}_{j, \setminus j} \mathbf{R}_{\setminus j, j} + \lambda \|\mathbf{R}_{\setminus j, j}\|_1 \right\}$$
 - 7: Set $\hat{G}^\lambda = (V, \hat{E}^\lambda)$ where $\hat{E}_{jk}^\lambda = \mathbb{I}\{\hat{\mathbf{R}}_{jk} \neq 0\}$
 - 8: **return** $\hat{G}^\lambda$
-

2.6. Stability selection

Using algorithm 2, we estimate the graph by selecting edges that indicate strong conditional dependence between two nodes. In this section, we employ a stability-based method to enhance the estimation performance. Instead of based on only one data realization, we aim to select those edges that *always* indicate strong conditional dependence under multiple realizations.

The stability selection method was first proposed by Meinshausen and Bühlmann (2010), and has become very successful in variable selection and structure estimation (Meinshausen and Bühlmann 2010, Liu *et al.* 2010, Shah and Samworth 2013). Using this method, the estimated graphs are more stable, in the sense that another set of realized data under the same data generating process would produce a graph that is very similar to the current one. Empirical results show that the stability selection method improves the selection accuracy.

Specifically, we first subsample without replacement from the original data, and then apply algorithm 2 for edge selection to each subsample. Next, we calculate the frequency that each edge is selected by the subsamples, and choose those with the highest frequencies to be selected. Suppose the size of each subsample is b .

There are in total $\binom{n}{b}$ subsamples. Ideally, we should use all subsamples to get the best result. However, Politis et al. (1999) showed that it suffices to use a large number L of

subsamples at random. For each subsample ℓ we estimate the sparse precision matrix. We then count the frequency that each edge is added to the graph using all L subsamples. Define

$$\hat{p}_{jk}^{(\lambda)} = \frac{\sum_{\ell=1}^L \mathbb{I}\{\hat{\mathbf{R}}_{jk}^{(\ell)} \neq 0\}}{L}.$$

where $\hat{\mathbf{R}}^{(\ell)}$ is obtained by the ℓ th subsample using algorithm (2). Taking different tuning parameter λ , we get a sequence of frequencies for each pair (j, k) . We can plot these frequencies against λ ; the resulted graph is called the stability path (figure 3). Stability selection chooses those edges that have high probability to be selected:

$$\hat{E}_{jk}^{(\lambda)} = \mathbb{I}\{\hat{p}_{jk}^{(\lambda)} > \alpha\}.$$

In practice, the estimated graph is insensitive to the choice of α (Liu *et al.* 2010, Meinshausen and Bühlmann 2010). The convention of $\alpha = 0.05$ is often chosen. For our purpose of independent variable selection, it suffices to choose $\alpha = 0$. The detailed procedure is presented in algorithm 3, which is a simple version of the stability selection. More complex methods can be considered. For example, instead of subsampling across time horizon, we can take the variable dimension, i.e. for each subsample, we take $(X_{1,J}, \dots, X_{n,J})$ where $J \subset \{1, 2, \dots, d\}$. Furthermore, we can employ a hybrid method, by taking subsamples from both the time and variable dimensions.

Algorithm 3

Stability Selection

-
- 1: $\hat{\mathcal{G}}^\lambda \leftarrow \text{Stability-Selection}(\mathbf{X}_1, \dots, \mathbf{X}_n; \lambda)$
 - 2: **input:** Historical return of d stocks over n periods $\mathbf{X}_1, \dots, \mathbf{X}_n \in \mathbb{R}^d$, tuning parameter λ , number of subsamples L , subsample size b and thresholding parameter α
 - 3: **output:** Estimated graph $\hat{\mathcal{G}}^\lambda$ corresponding to the tuning parameter λ
 - 4: **for** $\ell \in 1, \dots, L$ **do**
 - 5: Randomly draw b samples from the original data, denoted by $\mathbf{X}_{i_1}, \dots, \mathbf{X}_{i_b}$, where $1 \leq i_1 < i_2 < \dots < i_b \leq n$
 - 6: Calculate Kendall's tau correlation matrix $\hat{\Sigma} \in \mathbb{R}^{b \times b}$, where $\hat{\Sigma}_{jk} = \sin\left(\frac{\pi}{2} \hat{\tau}_{jk}(\mathbf{X}_{i_1}, \dots, \mathbf{X}_{i_b})\right)$ and $\hat{\tau}_{jk}(\mathbf{X}_{i_1}, \dots, \mathbf{X}_{i_b}) := \frac{2}{b(b-1)} \sum_{1 \leq \ell < \ell' \leq b} \text{sign}[(\mathbf{X}_{i_\ell j} - \mathbf{X}_{i_{\ell' j}}) \cdot (\mathbf{X}_{i_\ell k} - \mathbf{X}_{i_{\ell' k}})]$
 - 7: Let $\tilde{\Sigma} \leftarrow \text{Proximal-Gradient}(\hat{\Sigma})$ as in algorithm 1
 - 8: Obtain $\hat{\mathbf{R}}$ by solving
$$\hat{\mathbf{R}} = \underset{\mathbf{R} \in \mathbb{R}^{d \times d}, \text{diag}(\mathbf{R}) = \mathbf{0}}{\text{argmin}} \sum_{j=1}^d \left\{ \mathbf{R}_{\setminus j, j}^T \tilde{\Sigma}_{\setminus j, \setminus j} \mathbf{R}_{\setminus j, j} - 2 \tilde{\Sigma}_{j, \setminus j} \mathbf{R}_{\setminus j, j} + \lambda \|\mathbf{R}_{\setminus j, j}\|_1 \right\}$$
 - 9: Set $s_{jk}^{(\lambda)} = s_{jk}^{(\lambda)} + \mathbb{I}\{\hat{\mathbf{R}}_{jk} \neq 0\}$.

```

10:   Let  $\hat{p}_{jk}^{(\lambda)} \leftarrow s_{jk}^{(\lambda)} / L$ 
11:   Obtain  $\hat{G}^\lambda = (V, \hat{E}^\lambda)$  where  $\hat{E}_{jk}^\lambda = \mathbb{I}\{\hat{p}_{jk}^\lambda > 0\}$ 
12:   return  $\hat{G}^\lambda$ 

```

2.7. Independent variable selection

We are now ready to present the complete algorithm for stock selection. Using stability selection, we obtain a stability path corresponding to a sequence of tuning parameters λ . We start from the smallest λ , so that the estimated graph is+ the most dense. We next increase λ , so that the estimated graph gets sparser as λ increases. We scan this stability path until we obtain a graph in which there are 10 isolated nodes. Figure 3 shows the stability path and selection procedure of the example of health care stocks discussed in section 4.1; figure 4 shows the evolvement of the estimated graph when scanning the stability path with increasing λ .

Algorithm 4 lists the primary steps: (1) calculating Kendall's tau statistic, (2) smoothly projected neighbourhood pursuit, (3) stability selection and (4) independent node selection. The input requires a sequence of λ that is dense enough. To improve efficiency, we employ the binary search method, as the number of isolated points is monotonically nondecreasing as λ increases.

Algorithm 4

Stock Selection

```

1:  $\hat{\mathcal{J}} \leftarrow \text{Stock-Selection}(X_1, \dots, X_m; \lambda_1, \dots, \lambda_m; k)$ 
2: input: Historical return of  $d$  stocks over  $n$  periods  $X_1, \dots, X_n \in \mathbb{R}^d$ ; a sequence of tuning parameters  $\lambda_1 < \lambda_2 < \dots < \lambda_m$ ; desired number of stocks to be selected  $k$ 
3: output: An index set  $\hat{\mathcal{J}} \subset \{1, \dots, d\}$  with  $|\hat{\mathcal{J}}| = k$ 
4: for  $r \in 1, \dots, m$  do
5:    $\hat{G}^{\lambda_r} \leftarrow \text{Stability-Selection}(X_1, \dots, X_d; \lambda_r)$  as in Algorithm 3
6:    $N_r \leftarrow$  number of nodes that are not connected to other nodes in  $\hat{G}^{\lambda_r}$ 
7:   if  $N_r = k$  then break
8:    $\hat{\mathcal{J}} \leftarrow$  the indices of the nodes that are not connected to other nodes in  $\hat{G}^{\lambda_r}$ 
9: return  $\hat{\mathcal{J}}$ 

```

3. Theoretical guarantees

In this section, we provide theoretical guarantees for the stability selection procedure. Our main theorem establishes the statistical rate of convergence of the estimated index set $\hat{\mathcal{J}}$ to its population quantity $J := \mathbb{E}[\hat{\mathcal{J}}]$, where the expectation is taken over the randomness in subsampling. We define $p_{jk}^{(\lambda)} := \mathbb{E}[\hat{p}_{jk}^{(\lambda)}]$. Recall that n is the sample size, d is the total

number of stocks, m is the number of regularization parameters λ used in obtaining the graph path and b is the subsample size in the stability selection.

The following theorem shows that the estimated index set is consistent (converges to its population version):

Theorem 3.1

Suppose that m is a polynomial of n , i.e. $m = n^c$ for some constant c , and

$|p_{jk}^{(\lambda)} - p_{j'k'}^{(\lambda)}| \geq 2 \sqrt{\frac{c' b \log(n \vee d)}{n}}$ for any pairs (j, k) and (j', k') and λ , where c' is a universal constant. Then we have

$$\mathbb{P}(\hat{J}=J) \rightarrow 1$$

as $n \rightarrow \infty$.

The second condition in Theorem 3.1 is satisfied if $b \log(n \vee d)/n \rightarrow 0$ and for large enough n . Therefore, our algorithm scales to high-dimensional models, where d could potentially grow exponentially with n (appendix 1).

4. Application to equity selection

As mentioned, to harvest rebalancing gains, our investment goal is to select stocks to be as independent of each other as possible. We apply the above graphical modelling algorithm 4 to select a portfolio of stocks.

Herein, we consider the problem of portfolio management. Standard portfolio theory advocates using a variety of stocks to diversify the risk. Instead of using all available stocks, we choose a few stocks that are as independent as possible. We then construct our portfolio only using the chosen stocks. This dynamic portfolio strategy is expected to perform better than the traditional Markowitz model implementing the buy-and-hold strategy in terms of return, Sharpe ratio and other performance measures.

4.1. Healthcare data

First, to illustrate the approach, we select data from the health care industry (50 stocks). The price and return data start from January 1992 and end at April 2014. As many of the stocks did not exist until recently, we truncate the data before 1998 and only look at stocks that existed on January 2 1998. That gives us 43 stocks to depict the power of our method in terms of harvesting rebalancing gains. To address missing data, we use the k-nearest neighbour method to impute the data. This can be implemented by the function ‘impute.knn’ in the package ‘impute’ using the statistical analysis software R.

We compare three portfolio strategies. The first one is a rebalancing strategy, in which the stocks are selected using the graphical model and stability selection method described in section 2. The second one is also a rebalancing strategy, but the stocks are selected using a factor model approach, which combines the principal component analysis and clustering

technique. The comparison between these two approaches demonstrates the effectiveness of independent stock selection. The third one is the buy-and-hold Markowitz portfolio strategy. Comparison between the rebalancing strategies and the Markowitz approach demonstrates the benefits of rebalancing.

4.1.1. Graphical model strategy—The practical implementation of graphical model strategy is detailed as follows. At the beginning of the investment horizon, we select k stocks to put in the rebalancing portfolio using the graphical model and stability selection approach. The stocks are selected using previous years' returns as training data, and are expected to be as independent to each other as possible in terms of the graphical metric introduced in section 2. We invest only in these stocks for the whole year, and the investment amount in each stock is rebalanced at some frequency (usually monthly, weekly or daily) to remain at the same weights as their initial values. For simplicity, we use equal weight for all stocks. The set of selected stocks is updated after one year, again based on a new training set of historical data, until the investment period ends.

Note that equal weight is a first benchmark for achieving rebalancing gains, with better performance than the traditional buy-and-hold under many time periods (Mulvey and Kim 2009). Improved performance can be obtained by rendering robust estimates of the expected values of each stock, but this extension requires considerably more information.

The graphical model stock selection procedure is described in the following steps. First, obtain the historical returns of all candidate stocks. Then compute the graph path of all the stocks for a sequence of regularization parameters $\lambda_1 < \dots < \lambda_m$ via the graphical model and stability selection method as described in algorithms 2 and 3. From the graph path, select one graph whose number of isolated nodes is equal to k . Specifically, one can start from λ_1 , when the graph is densest (e.g. see first graph in figure 4). Usually, all the nodes are connected, so we proceed to count the number of isolated nodes of the next graphs, until this number is greater than or equal to k . The isolated nodes on the final graph are the stocks we select to put in the rebalancing portfolio.

In our real data experiment of the health care stocks, we consider two scenarios: one is based on monthly return data and the rebalancing frequency is monthly, the other is based on daily data the rebalancing frequency is daily. We use previous four year's stock returns as the training data, and select 10 among the total 43 stocks to put in the rebalancing portfolio (i.e. $k = 10$). The historical returns are stored in a 48×43 data matrix for monthly returns, or 1008×43 for daily returns. We then compute the graph path using the data matrix and algorithms 2 and 3. In practice, the procedures can be conveniently implemented by the functions in the High-dimensional Undirected Graph Estimation ('huge') package (Zhao, Liu *et al.* 2012) using the statistical analysis software R. We briefly describe the usage of this package and the parameter settings in our example.

First, input the data matrix to the 'huge' function, and select 'mb' for the 'method' option. This function computes a list of Markovian graphs corresponding to a sequence of regularization parameters, and the option 'mb' stipulates that the smoothly projected neighbourhood selection method (algorithm 2) is adopted for graph estimation. The

sequence of regularization parameters $\lambda_1 < \dots < \lambda_m$ is automatically generated by the function, where the two ends λ_1, λ_m corresponds to the densest and sparsest (empty) graphs of the 48 stock returns. The output of the function 'huge' is an R object of the class 'huge', which is then input into the function 'huge.select' for implementing the stability selection. In this function, subsampling is performed to select edges and obtain a graph path as described in algorithm 3, for the sequence of regularization parameters chosen before. An example of the graph path is shown in figure 3. There are several arguments that controls the subsampling. The 'stars.subsample.ratio' is the ratio between the sizes of the subsample (denoted by b in algorithm 3) and the entire sample (n). The default is $10\sqrt{n}/n$ if $n > 144$ and 0.8 if $n < 144$. Another argument is 'rep.num', which is the number of subsamples (denoted by L in algorithm 3), which we take to be 50. Lastly, scan the graph path to obtain the graph with 10 isolated nodes. Sometimes this number might not be able to achieved if the sequence of regularization parameters is not refined enough, and we choose the closest number. In the function 'huge', we can stipulate the refinement of this sequence. Specifically, 'nlambda' controls the number of regularization parameters m in the sequence.

In practice, $m = 30$ often results in 9–11 stocks, and with $m = 50$, the number of selected stocks can be exactly 10.

4.1.2. PCA clustering strategy—The second strategy is also based on rebalancing, but the set of k stocks is selected using a factor model. Similar to the previous strategy, we use historical returns as our training data for stock selection. We project the stocks into the space spanned by the leading three principle components (factors), and cluster their low dimension projections into k groups. The intuition is that under the factor model, stock returns are mostly determined by the leading three factors. Hence those stocks belonging to different clusters should have low correlations as their factor loadings are rather different. We select the stocks closest to each of the k cluster centres and use the k stocks to construct our portfolio.

Specifically, we take the data matrix containing the stock returns, and perform the eigen decomposition of its empirical covariance matrix. In particular, the eigenvectors are orthonormal. We take the three leading eigenvectors (that is, eigenvectors corresponding to the three largest eigenvalues). Each dimension represents one stock, so the j -th coordinates of the three vectors are the loadings (weights) of the three principal components corresponding to the j -th stock. We cluster the stocks into k groups by their factor loadings using the k -means algorithm, where the L_2 -norm of the factor loadings is used as the distance metric for the stocks. The k -means algorithm is described as follows. Given an initial clustering of k groups, compute the mean of each group as the new cluster centre. Then calculate the distance of each member to the k centres, and cluster the members into k new groups according to their closest centres. Next calculate the mean of the each group to update the cluster centres. The process is repeated until convergence. Note that in our example, the members are the three-dimensional factor loadings of each stock. In the R software, we use the build-in function 'kmeans' to implement the algorithm. The inputs are the 48×3 factor loading matrix, and the parameter k to specify the number of clusters. We take $k = 10$. The output is an R object containing the information of the cluster centres.

Finally, we select the ones closest to each cluster centre as the 10 stocks to be invested in the rebalancing portfolio.

4.1.3. Markowitz strategy—The last strategy we consider is Markowitz portfolio selection. This traditional selection method is characterized by its ‘buy-and-hold’ nature, which ignores the rebalancing gains. Stocks are selected once every year and are held for the entire year. The description of Markowitz strategy can be found in multiple literature, e.g. section 6.6 of Luenberger (2013).

The comparison results appear in tables 2 and 3. The results show that the selection method based on graphical model is better in return, Sharpe ratio and max drawdown. The ‘buy-and-hold’ Markowitz strategy produces the least favourable performance, with low returns, high standard deviations and large maximum drawdown.

4.1.4. Transaction cost—We discuss the impact of transaction cost on the portfolio returns. For the rebalancing strategy, there are two sources of transaction costs: one is from daily (monthly) rebalancing, and the other is from adjusting the set of stocks selected at the beginning of each year. For the Markowitz strategy, transaction cost only occurs at the beginning of each year for stock selection. In tables 4 and 5, we calculate the annualized turnovers and the impact on the portfolio returns under different rates of transaction costs. In particular, the first column is the annualized turnovers of the three strategies, and the last three columns are the adjusted annualized returns after accounting for transaction costs, at the rates of 0.1, 0.5 and 1%.

An investor can thereby see the impact of these costs on performance depending upon their circumstances. In the experiment result, we can see that the rebalancing strategy indeed requires significantly more trading amounts than the buy-and-hold Markowitz strategy. However, the impact of transaction cost on the portfolio return is relatively small for 1% transaction cost.

4.2. Large stock data

In our second real data experiment, we apply our graphical model strategy to a large data-set including 1500 stocks over a 34-year investment horizon from 1981–2014. In the first step, we implement a momentum-based investment procedure (Asness *et al.* 2013) to preselect 50 stocks. At selected time junctures, we rank stocks in the S&P 1500 index according to their performance over a predetermined back-test period (typically 3–12 months). Under momentum, stocks with the best historical performance have a higher probability of better performance over the subsequent holding period. For our purposes, we take the top 50 stocks out of 1500 as a target list based on a six-month look-back. In the second step, we compare four strategies for choosing a portfolio among the 50 stocks: (1) rebalancing with graphical model selection; (2) rebalancing with PCA clustering selection; (3) traditional Markowitz model; and (4) rebalancing with all 50 stocks. Under alternatives 1, 2 and 4, we apply an equal weighted mix—either 10 stocks for 1 and 2, or 50 stocks for 4.

Herein, the ranking and independent stock selection are conducted annually, followed by daily rebalancing until the next ranking and selection date. We use two-year historical

returns as the training data-set to compute the independence set. For stability selection, we take $L = 50$ and $m = 30$. Our method exhibits significant positive results in the following aspects:

1. Our method substantially outperforms the classical Markowitz strategy. This shows that our approach of applying graphical model combined with rebalancing strategy is able to reap rebalancing gains compared with the traditional ‘buy-and-hold’ strategy. This empirically conforms with the theories of rebalancing gains discussed in Luenberger (2013) and Mulvey *et al.* 2001, etc.
2. Compared with rebalancing strategy using all 50 stocks, the rebalancing portfolio with 10 stocks selected by the graphical model has higher returns, better Sharpe ratios and smaller max-drawdowns. This not only confirms the advantage of using independent stocks in such rebalancing strategy, but also shows that the stocks selected by our graphical method indeed exhibit less dependence than the entire pool of stocks.
3. Compared with the factor-based clustering method, the approach based on graphical model has better performance for the momentum-based investment procedure (figure 5, table 6. Note that the y-axis of figure 5 is log-scale). This suggests that our latent hierarchical framework is able to effectively capture the inter-dependency structure among stocks, while the factor model only partially captures such relationships.

4.2.1. Transaction cost—Similar to the previous example, we calculate the annualized turnovers and impact of transaction costs, which are summarized in table 7.

4.2.2. Computation time—Our stability based stock selection algorithm potentially requires many iterations, due to subsampling in the stability path. However, as computation of each iteration is very fast, the overall procedure performs reasonable in time. In the health care example, for the monthly return data, we use a four-year sliding window as training set to select independent stocks. So $n = 48$ and $d = 43$. As mentioned above, the number of subsamples L is 50 and the number of tuning parameters m is 30. The computation time is 15.308 s.

With daily data, we have $n = 1008$ and $d = 42$. With the same L and m , the computation time is 15.680 s. Substantially, the computation time is relatively invariant to the sample size (n). This is because the data size is only involved in the computation of the Kendall’s tau statistic, which costs little time. The rest of the algorithm is not related to the sample size.

Lastly, in the large-scale stock example, we take a two-year sliding window, so $n = 504$ and $d = 50$. Same as the above example, we take $L = 50$ and $m = 30$. The computation time is 16.616 s.

All the results are run on a Macintosh OS X with 2.7 GHZ Intel Core i5 CPU and 16GB 1600 MHZ DDR3 Memory.

5. Independence measure

In previous sections, we presented our stock selection method based on graphical model approach to construct a portfolio that exploits rebalancing gains. The crucial idea is to select stocks that exhibit the highest level of independence. Empirical results show that our method for independent stock selection is able to harvest such rebalancing gains. In this section, we propose a new measure that characterizes the level of independence between stocks. To this end, we apply a novel distribution-free independence test. The level of independence is characterized by the p -value of this test. The stocks selected by graphical model method indeed exhibit higher level of independence compared with the methods of PCA selection, random selection or selecting all available stocks, which justifies the validity of our method.

Formally, this test is formulated as follows: let $\mathbf{X} = (X_1, \dots, X_d)$ be a d -dimensional random vector. In our example, $\mathbf{X}$ is the return of d selected stocks. We aim to test the null hypothesis:

$$\mathbf{H}_0: X_1, X_2, \dots, X_d \text{ are mutually independent.} \quad (5.1)$$

This problem has been well studied in statistics literature under the Gaussian assumption, for example Anderson (1958), Roy and Roy (1957), Nagao (1973) in low dimensional settings and Bai *et al.* (2009), Cai and Ma (2013), Jiang *et al.* (2013), Johnstone (2001), Ledoit and Wolf (2002), Schott (2005), Srivastava (2005) in high dimensions. Moreover, Jiang (2004) and Zhou (2007) relaxed the Gaussian assumption to distributions with high order moments.

We aim to test (5.1) based on the rank coefficient statistics. Our method is state-of-art, as it does not require any distributional assumption on $\mathbf{X}$, and works well under high-dimensional settings where d is allowed to grow at the rate of $\log d = \alpha(n^{1/3})$.

Given realizations $\mathbf{X}_1, \dots, \mathbf{X}_n$, the test statistic is defined as follows:

$$\hat{T}(\mathbf{X}_1, \dots, \mathbf{X}_n) := \max_{j < k} |\hat{\tau}_{jk}(\mathbf{X}_1, \dots, \mathbf{X}_n) - \mathbb{E}_0[\hat{\tau}_{jk}(\mathbf{X}_1, \dots, \mathbf{X}_n)]|,$$

where $\hat{\tau}_{jk}(\mathbf{X}_1, \dots, \mathbf{X}_n)$ is the Kendall's tau statistic defined by (2.5) and $\mathbb{E}_0$ is the expectation taken under $\mathbf{H}_0$. Hence $\mathbb{E}_0[\hat{\tau}_{jk}]$ is the population version of Kendall's tau rank correlation under the null hypothesis. It can be shown under $\mathbf{H}_0$ as defined in (5.1), $\mathbb{E}_0[\hat{\tau}_{jk}] = 0$. The following theorem shows the asymptotic behaviour of the test statistic $\hat{T}$ under $\mathbf{H}_0$.

Theorem 5.1

Under the scaling that $\log d = \alpha(n^{1/3})$ and the null hypothesis $\mathbf{H}_0$, we have for any $x \in \mathbb{R}$

$$|\mathbb{P}(\hat{T}^2/\sigma_0^2 - 4 \log d + \log \log d \leq x) - \exp(-\exp(-x/2)/\sqrt{8\pi})| = o(1),$$

where σ_0^2 is the variance of $\hat{\tau}$ under the $\mathbf{H}_0$, and when $\mathbf{H}_0$ is as in (5.1), $\sigma_0^2 = \frac{2(2n+5)}{9n(n-1)}$.

Theorem 5.1 shows that the distribution function of $\hat{T}^2/\sigma_0^2 - 4 \log d + \log \log d$ can be approximated by a Gumbel distribution when n is large. The proof can be found in the proof of Theorem 3.2 in Han and Liu (2014). Based on this theorem, we construct the test statistic as follows:

$$\phi(\mathbf{X}_1, \dots, \mathbf{X}_n) = \mathbb{I}\{\hat{T}^2 > \sigma^2(4 \log d - \log \log d - 2 \log \sqrt{8\pi} - 2 \log \log(1 - \alpha)^{-1})\}.$$

The validity of this test ϕ follows immediately from Theorem 5.1, as the type I error is approximately α when n is large:

$$\mathbb{P}(\phi(\mathbf{X}_1, \dots, \mathbf{X}_n) = 1 | \mathbf{H}_0) = \alpha + o(1).$$

We also provide power analysis that controls the type II error of the test ϕ . The following theorem shows that the power goes to 1 when n goes to infinity. We consider the following family of alternative hypotheses that are defined in terms of the population Kendal' tau matrix: $\mathcal{H}_1 = \{\mathcal{T} \in R(a)\}$, where $\mathcal{T} = [\mathbb{E}[\hat{\tau}_{jk}]]_{1 \leq j, k \leq d} \in \mathbb{R}^{d \times d}$, and

$$R(a) := \left\{ \mathbf{M} \in \mathbb{R}^{d \times d} : \text{diag}(\mathbf{M}) = \mathbf{I}_d, \mathbf{M} = \mathbf{M}^T, \max_{1 \leq j < k \leq d} |\mathbf{M}_{jk}| \geq a \sqrt{\log d/n} \right\}.$$

Theorem 5.2

There exists some large enough constant $C > 0$ such that

$$\inf_{\mathcal{T} \in R(C)} \mathbb{P}(\phi = 1) = 1 - o(1).$$

See Theorem 4.2 of Han and Liu (2014) for the proof.

This test is applied to the stock selection problem. We measure the independence level of stocks selected by different methods for the healthcare portfolio. Specifically, we compare the four methods: (1) graphical model approach; (2) PCA dimension reduction followed by clustering; (3) random selection; (4) selecting all stocks. The p -values demonstrate the independence level of the selected stocks. For clearance of presentation we rescale the p -values by first taking the logarithm and then adding a constant to make them positive. The rescaled p -values are shown in figure 6 for yearly selection and in figure 7 for monthly selection of the health care stock data example in section 4.1. From the figures, we can see that the p -values of stocks selected by graphical model approach are higher than other methods, indicating that our method indeed is able to select stocks that possess more independence. Applying this measure for large-scale stock example in section 4.2, the result is shown in figure 8.

Also, notice that in all of the examples, the independence level of stocks decreases drastically during 2008/2009 crash periods under the proposed independence measure, which agrees with theoretical and empirical evidence that assets are more correlated in crash

periods than in normal periods. A more interesting fact revealed by all three figures appears to be that there is a permanent change in the correlation structure among financial assets before and after the crash. For example, during the period 2010–2014, the correlation level stayed high despite the rebounding market, as opposed to the period 2002–2006 when the market was also growing. This may be explained by the fact that investors have become more cautious and behaved qualitatively different after the financial crisis.

6. Conclusion

Portfolio performance can be enhanced by means of rebalancing gains. Theoretically, when the assets are independent and transaction costs are small, we are able to harvest the rebalancing gains compared with a traditional buy-and-hold strategy. However, implementation of this strategy is difficult as stocks possess substantial co-movements. In this paper, we considered the challenging problem of selecting stocks that are ‘as independent as possible’, and applied our method to implement the rebalancing tactic. The method uses the ideas of elliptical-copula graphical models. Under such a modelling strategy, the notion of independence is defined by the conditional independence relationships among the latent Gaussian random variables that underlie the observed variables. We propose an algorithm via ℓ_1 -penalization and smoothed projection methods for estimating the latent graphical structure. This algorithm not only has a fast convergence rate, but also scales to potentially a very large number of stocks. We also apply stability selection to improve the estimation performance. The stocks to add to the portfolio are selected by a two-step thresholding method.

Empirically, the portfolio strategy implementing the graphical model selection significantly outperforms the traditional buy-and-hold strategy such as the Markowitz strategy, as well as alternative rebalancing strategies such as PCA dimension reduction combined with clustering. This not only empirically shows the advantages of investing independent stocks in the rebalancing strategy, but also depicts the power of latent Markovian graphical methods in modelling the dependence structure among financial assets.

Acknowledgments

Funding

This research is partially supported by [NSF CAREER Award DMS1454377], [NSF IIS1408910], [NSF IIS1332109], [NIH R01MH102339], [NIH R01GM083084], [NIH R01HG06841], and SandPointe, LLC.

References

- Anderson, TW. An Introduction to Multivariate Statistical Analysis. New York: Wiley; 1958.
- Asness CS, Moskowitz TJ, Pedersen LH. Value and momentum everywhere. *J. Finance*. 2013; 68:929–985.
- Bai Z, Jiang D, Yao JF, Zheng S. Corrections to IRT on large-dimensional covariance matrix by RMT. *Ann. Stat.* 2009; 37:3822–3840.
- Cai TT, Ma Z. Optimal hypothesis testing for high dimensional covariance matrices. *Bernoulli*. 2013; 19:2359–2388.
- Cherubini, U., Luciano, E., Vecchiato, W. *Copula Methods in Finance*. Chichester: John Wiley & Sons; 2004.

- Costa Dias, AD. Doctoral Thesis ETH No. 15283. ETH Zurich; 2004. Copula inference for finance and insurance.
- Dempster AP. Covariance selection. *Biometrics*. 1972; 28:157–175.
- Dempster MAH, Evstigneev IV, Schenk-Hoppé KR. Financial markets: The joy of volatility. *Quant. Finance*. 2008; 8:1–3.
- Dempster MAH, Evstigneev IV, Schenk-Hoppé KR. Growing wealth with fixed-mix strategies. The Kelly Capital Growth Investment Criterion: Theory and Practice. 2010; 3:427–455.
- Dias, A., Embrechts, P. Manuscript. ETH Zurich; 2004. Dynamic copula models for multivariate high-frequency data in finance.
- Embrechts, P., Lindskog, F., McNeil, A. Modelling dependence with copulas and applications to risk management. In: Rachev, S., editor. *Handbook of Heavy Tailed Distributions in Finance*. Vol. 8. Amsterdam: Elsevier; 2003. p. 329–384.
- Friedman J, Hastie T, Höfling H, Tibshirani R. Pathwise coordinate optimization. *Ann. Appl. Stat.* 2007; 1:302–332.
- Gorton GB, Hayashi F, Rouwenhorst KG. The fundamentals of commodity futures returns. *Rev. Finance*. 2013; 17:35–105.
- Gorton, G., Rouwenhorst, KG. Technical report. National Bureau of Economic Research; 2004. Facts and fantasies about commodity futures.
- Han F, Liu H. Transelliptical component analysis. *Advances in Neural Information Processing Systems*. 2012; 25:171–179.
- Han, F., Liu, H. Distribution-free tests of independence with applications to testing more complex structures. Technical report. 2014.
- Higgs H, Worthington A. Stochastic price modeling of high volatility, mean-reverting, spike-prone commodities: The Australian wholesale spot electricity market. *Eng. Econ*. 2008; 30:3172–3185.
- Jiang T. The asymptotic distributions of the largest entries of sample correlation matrices. *Ann. Appl. Probab.* 2004; 14:865–880.
- Jiang T, Yang F, et al. Central limit theorems for classical likelihood ratio tests for high-dimensional normal distributions. *Ann. Stat.* 2013; 41:2029–2074.
- Johnstone IM. On the distribution of the largest eigenvalue in principal components analysis. *Ann. Stat.* 2001; 29:295–327.
- Ledoit O, Wolf M. Some hypothesis tests for the covariance matrix when the dimension is large compared to the sample size. *Ann. Stat.* 2002; 30:1081–1102.
- Li DX. On default correlation: A copula function approach. *J. Fixed Income*. 2000; 9:43–54.
- Lindskog, F., McNeil, AJ., Schmock, U. Kendall's tau for elliptical distributions. In: Bol, G.Nakhaeizadeh, G.Rachev, ST.Ridder, T., Vollmer, K-H., editors. *Credit Risk-Measurement, Evaluation and Management*. Heidelberg: Physica—Verlag; 2003. p. 149–156.
- Liu H, Han F, Yuan M, Lafferty J, Wasserman L. High dimensional semiparametric gaussian copula graphical models. *Ann. Stat.* 2012a; 40:2293–2326.
- Liu H, Han F, Zhang CH. Transelliptical graphical models. *Advances in Neural Information Processing Systems*. 2012b; 25:809–817.
- Liu, H., Han, F., Zhang, CH. Transelliptical graphical modeling under a hierarchical latent variable framework. Technical report. 2013.
- Liu H, Roeder K, Wasserman L. Stability approach to regularization selection (stars) for high dimensional graphical models. *Advances in Neural Information Processing Systems*. 2010; 23
- Luenberger, DG. *Investment Science*. 2nd. New York: Oxford University Press; 2013.
- Meinshausen N, Bühlmann P. High-dimensional graphs and variable selection with the lasso. *Ann. Stat.* 2006; 34:1436–1462.
- Meinshausen N, Bühlmann P. Stability selection. *J. Roy. Stat. Soc.: Ser. B (Stat. Method.)*. 2010; 72:417–473.
- Mulvey JM. Long-short versus long-only commodity funds. *Quant. Finance*. 2012; 12:1779–1785.
- Mulvey JM, Kaul SSN, Simsek KD. Evaluating a trend-following commodity index for multi-period asset allocation. *J. Altern. Investments*. 2004; 7:54–69.

- Mulvey, JM., Kim, WC. *Encyclopedia of Quantitative Finance*. Chichester: John Wiley and Sons; 2009. Constantly rebalanced portfolio— Is mean reversion necessary?.
- Mulvey JM, Lu N, Sweemer J. Rebalancing strategies for multi-period asset allocation. *J. Wealth Manage.* 2001; 4:51–58.
- Mulvey JM, Ural C, Zhang Z. Improving performance for long-term investors: Wide diversification, leverage, and overlay strategies. *Quant. Finance.* 2007; 7:175–187.
- Nagao H. On some test criteria for covariance matrix. *Ann. Stat.* 1973; 1:700–709.
- Nesterov, Y., Nesterov, IE. *Introductory Lectures on Convex Optimization: A Basic Course*. Vol. 87. Springer; 2004.
- Politis, D., Romano, J., Wolf, M. *Subsampling*, Springer series in statistics. New York: Springer; 1999.
- Roy, SN., Roy, S. *Some Aspects of Multivariate Analysis*. New York: Wiley; 1957.
- Schott JR. Testing for complete independence in high dimensions. *Biometrika.* 2005; 92:951–956.
- Shah RD, Samworth RJ. Variable selection with error control: Another look at stability selection. *J. Roy. Stat. Soc.: Ser. B (Stat. Methodol.)*. 2013; 75:55–80.
- Srivastava MS. Some tests concerning the covariance matrix in high dimensional data. *J. Jpn. Stat. Soc.* 2005; 35:251–272.
- Vanderbei RJ. Loqo: An interior point code for quadratic programming. *Optim. Methods Softw.* 1999; 11:451–484.
- Xue L, Zou H. Regularized rank-based estimation of high-dimensional nonparanormal graphical models. *Ann. Stat.* 2012; 40:2541–2571.
- Zhao T, Liu H, Roeder K, Lafferty J, Wasserman L. The huge package for high-dimensional undirected graph estimation in R. *J. Mach. Learn. Res.* 2012; 13:1059–1062. [PubMed: 26834510]
- Zhao T, Roeder K, Liu H. Smooth-projected neighborhood pursuit for high-dimensional nonparanormal graph estimation. *Advances in Neural Information Processing Systems.* 2012; 25:162–170.
- Zhou W. Asymptotic distribution of the largest off-diagonal entry of correlation matrices. *Trans. Am. Math. Soc.* 2007; 359:5345–5363.

Appendix 1

Proof of theorem 3.1

We first prove the convergence of $\hat{p}_{jk}^{(\lambda)}$ for each j, k and λ . Note that $\hat{p}_{jk}^{(\lambda)}$ is a U-statistic of order b . Then using Hoeffding's inequality, we have for any $t > 0$,

$$\mathbb{P}(|\hat{p}_{jk}^{(\lambda)} - p_{jk}^{(\lambda)}| > t) \leq 2 \exp \left(-2nt^2/b \right).$$

By union bound, we have

$$\mathbb{P} \left(\exists r, j, k \text{ such that } |\hat{p}_{jk}^{(\lambda_r)} - p_{jk}^{(\lambda_r)}| > t \right) \leq md(d-1) \exp \left(-2nt^2/b \right) \leq \exp(\log m + 2 \log d - 2nt^2/b).$$

For any $\delta \in (0, 1)$, we let t satisfy $\delta = \exp(\log m + 2 \log d - 2nt^2/b)$ in the above inequality. Hence, we get with probability $1 - \delta$,

$$|\hat{p}_{jk}^{(\lambda_r)} - p_{jk}^{(\lambda)}| \leq \sqrt{\frac{b(\log m + 2 \log d - \log \delta)}{2n}}, \quad (\text{A1})$$

for all $r = 1, \dots, m$, and j, k such that $1 \leq j < k \leq d$.

Define the event $\mathcal{E} = \left\{ |\hat{p}_{jk}^{(\lambda_r)} - p_{jk}^{(\lambda_r)}| \leq \sqrt{\frac{c' b \log(n \vee d)}{n}} \text{ for any } j, k, r \right\}$, where c' is some universal constant.

Letting $\delta = (n \vee d)^{-1}$ in (A1) and by the fact that $m = n^c$, we have $\mathbb{P}(\mathcal{E}^c) \leq (n \vee d)^{-1}$. On the

event $\mathcal{E}$, and by the condition that $|p_{jk}^{(\lambda)} - p_{j'k'}^{(\lambda)}| \geq 2 \sqrt{\frac{c' b \log(n \vee d)}{n}}$, we have that for any λ_r ,

the order of $\hat{p}_{jk}^{(\lambda_r)}$ is the same as $p_{jk}^{(\lambda_r)}$ for $1 \leq j < k \leq d$. Thus, on the event $\mathcal{E}$, we have

$\hat{\mathcal{G}}^{(\lambda_r)} = \mathcal{G}_0^{(\lambda_r)}$, where $\mathcal{G}_0^{(\lambda_r)} = \mathbb{E}[\hat{\mathcal{G}}^{(\lambda_r)}]$, where the expectation is taken with respect to randomness in subsampling. Thus $\mathcal{E}$ implies the event that $\hat{\mathcal{J}} = \mathcal{J}$. Therefore,

$$\mathbb{P}(\hat{\mathcal{J}} \neq \mathcal{J}) \leq \mathbb{P}(\mathcal{E}^c) \rightarrow 0.$$

which completes the proof.

Figure 1.

Latent hierarchical framework for elliptical-copula graphical model. All three layers share the same graph structure, which is defined by $\mathcal{G}_{jk} = \mathbb{I}\{\Theta_{jk} \neq 0\}$, but have different interpretations.

Figure 2.
The soft thresholding function $\tilde{T}_{jk}$.

Figure 3.

Stability path and the selection procedure of the health care stock example. In this figure, a curve is the stability path of an edge. There are in total $d(d-1)/2$ paths. As we scan the stability paths with increasing λ , fewer edges are selected, and more isolated nodes emerge. The final λ is chosen at a point where the corresponding graph has exactly 10 isolated nodes (when $\lambda = 0.5006$ in this example). The red curves are paths of edges selected in the final graph.

Figure 4.

Evolution of the graph generated by stability selection with increasing λ . The initial graph is very dense. As we increase λ , the graph becomes sparser and more isolated nodes emerge. The final graph is selected when there are exactly 10 isolated nodes, when $\lambda = 0.5006$ in this example.

Figure 5.
Comparison of strategies with stocks based on value with yearly selection and daily rebalancing.

Figure 6.
Comparisons of independence level of selected stocks based on yearly selection in the health care stock example.

Figure 7.
Comparisons of independence level of selected stocks based on monthly selection for the health care stock example.

Figure 8.
Comparisons of independence level of selected stocks in the large-scale stock data example.

Table 1

Level of rebalancing gains for stylized example.

Volatility (%)	Rebalancing gains (%)
20	1.8
40	7.2
60	16.2

Table 2

Comparisons of stock strategies based on monthly data with monthly rebalancing.

	Ann. Return (%)	Ann. SD (%)	Sharpe	Max drawdown (%)
Graphical model	15.32	16.51	0.87	31.16
PCA clustering	14.08	20.28	0.64	36.47
Markowitz	4.23	32.03	0.10	66.89

Table 3

Comparisons of stock strategies based on daily data with daily rebalancing.

	Ann. Return (%)	Ann. SD (%)	Sharpe	Max drawdown (%)
Graphical model	20.34	21.09	0.92	41.70
PCA clustering	15.80	23.48	0.63	48.00
Markowitz	10.80	23.33	0.42	60.95

Table 4

Impact of transaction costs on portfolio returns based on monthly data.

	Ann. Turnover (%)	0.1% Trans. Cost	0.5% Trans. Cost	1% Trans. Cost
Graphical model	37.31	15.27	15.10	14.88
PCA clustering	31.30	14.04	13.90	13.72
Markowitz	22.70	4.20	4.11	3.99

Table 5

Impact of transaction costs on portfolio returns based on daily data.

	Ann. Turnover (%)	0.1% Trans. Cost	0.5% Trans. Cost	1% Trans. Cost
Graphical model	112.6	26.62	26.31	25.92
PCA clustering	101.1	15.68	15.20	14.56
Markowitz	3.08	10.80	10.78	10.77

Table 6

Comparison of strategies using momentum-based pre-selected stocks.

	Ann. Return (%)	Ann. SD (%)	Sharpe	Max drawdown (%)
Graphical model	31.00	22.70	1.32	67.44
PCA clustering	30.31	28.29	1.04	74.20
Markowitz	13.15	61.76	0.20	97.01
All 50 stocks	27.66	22.61	1.18	74.86
S&P500	7.65	18.35	0.36	57.60

Table 7

Impact of transaction costs on portfolio returns for large stock data example.

	Ann. Turnover (%)	0.1% Trans. Cost	0.5% Trans. Cost	1% Trans. Cost
Graphical Model	70.57	30.91	30.51	29.93
PCA Clustering	50.59	30.24	29.96	29.58
Markowitz	1.95	13.15	13.14	13.13
All 50 Stocks	167.0	27.44	26.39	24.39