Optimal trading strategies — a time series approach

Peter A. Bebbington[†] [‡]
[†]*Department of Physics and Astronomy, University College London, Gower Street, London WC1E 6BT, U.K. and*
[‡] *Centre for Doctoral Training in Financial Computing & Analytics,*
University College London, Malet Place, London WC1E 7JG, U.K.

Reimer Kühn*
**Department of Mathematics, King's College London, Strand, London WC2R 2LS, U.K.*

(Dated: JSTAT, Received 29 Sept 2015, accepted 29 Feb 2016, published 19 May 216)

Abstract

Motivated by recent advances in the spectral theory of auto-covariance matrices, we are led to revisit a reformulation of Markowitz' mean-variance portfolio optimization approach in the time domain. In its simplest incarnation it applies to a single traded asset and allows to find an optimal trading strategy which — for a given return — is minimally exposed to market price fluctuations. The model is initially investigated for a range of synthetic price processes, taken to be either second order stationary, or to exhibit second order stationary increments. Attention is paid to consequences of estimating auto-covariance matrices from small finite samples, and auto-covariance matrix cleaning strategies to mitigate against these are investigated. Finally we apply our framework to real world data.

Key words: risk measure and management, stochastic processes

I. INTRODUCTION

When seeking an optimal strategy for capital allocation one can adopt a dynamic programming approach that requires solving a Hamilton-Jacobi-Bellman or Bellman equation [1–8] to find such a strategy. An alternative approach, typically applied to single period problems, is mean-variance optimization, which forms the basis of Markowitz' portfolio optimization theory [9]. This approach has a rich history in economic research and industrial practice [10–14]. One of the main reasons for its popularity is clearly its conceptual simplicity, which helps building an intuition about the nature of risk and its relation to an investment's return.

The last couple of decades has seen many physicists becoming interested in this very same question [15–26]. Key issues addressed in these studies concern the effects that sampling noise is likely to have on the measurement of correlations or covariances in large portfolios, the way in which such sampling noise is going to affect the solution of a subsequent mean-variance portfolio optimization problem, and the design of methods to mitigate against adverse effects of such sampling noise.

The bedrock of most of these studies is the theory of random sample covariance matrices [27]. Their spectral theory was pioneered by Marčenko and Pastur [28] in the 1960's. It has indeed been observed that — apart from a number of *large* eigenvalues — the bulk of the spectrum of sample-covariance matrices of asset returns in various markets is very close to the form predicted by Marčenko and Pastur for sample covariance matrices of i.i.d. random data; see e.g. [15, 16]. This type of comparison between market data and a null-model defined by random data could then be used to devise theory-guided ways of distinguishing between information and noise in market data, and thereby to devise methods to clean covariance matrices of asset returns for the purpose of their subsequent use in portfolio optimization, with the effect of improving risk-return characteristics [15, 17–26].

The present study was triggered by the fact that the spectral theory of sample *auto-covariance* matrices — the analogue of [28] in the time domain — has recently become available [29]. This leads us to revisit the analogue of Markowitz mean-variance optimization in the time domain [30], which in its simplest incarnation allows to find an optimal trading strategy for a single traded asset over a finite (discrete) time horizon. We investigate this setup for a range of synthetic processes, taken to be either second order stationary, or to exhibit second order stationary increments, and we systematically study the effects of sampling noise on optimal strategies and on risk-return characteristics. Finally we apply our framework to daily returns of the S&P500 index, and we explore how results obtained for spectra of sample auto-covariance matrices obtained in [29] could then be used as a guide to clean sample auto-covariance matrices in a spirit analogous to that used for sample-covariance matrices in the context of portfolio optimization.

We note at the outset that we regard this as an exploratory study, and that we ignore economic factors such as discouning and agents' asymmetric perceptions of gains and losses in the present paper. We expect that the primary area of application of our techniques would be in the high-frequency domain, as return auto-correlations will be most prominent at short times. We note, however, that much of our analysis is about effects of sampling noise on optimal trading strategies, which is relevant at *all* time scales, and thus also for weakly correlated data.

The remainder of this paper is organized as follows. In Sect. II we briefly describe Markowitz' approach to portfolio optimization, and its translation into the time domain. In Sect. III we provide results for synthetic processes, and numerically investigate the influence of sampling noise on optimal strategies and risk-return profiles. In Sect. IV we look at optimal trading strategies for empirical data, using the S&P500 index as an example and we investigate the effect of auto-covariance matrix cleaning on risk-return profiles, based on comparing auto-covariance spectra for the S&P500 and expected spectra for a process with uncorrelated increments. Sect. V finally is devoted to a final overview, and an outlook on promising future research directions.

II. PORTFOLIO OPTIMISATION

A. The Markowitz Set-Up

In the simplest version of mean-variance portfolio optimization one considers a set of N tradable assets $i = 1, \dots, N$. It is usually assumed that these do not include complex financial instruments such as derivatives, options and futures. An investor can take positions on these assets. We will use π_i to denote the position on asset i, using the convention that $\pi_i > 0$ represents a long position (buying the asset), whereas $\pi_i < 0$ represents a short position (selling asset). With r_i denoting the (random) return on the i-th asset, the return on the entire portfolio with positions $\boldsymbol{\pi} = (\pi_1, \pi_2, \dots, \pi_N)'$ is given by

$$R(\boldsymbol{\pi}) = \sum_{i=1}^{N} \pi_i r_i = \boldsymbol{\pi}' \boldsymbol{r} \ , \tag{1}$$

where $\boldsymbol{r} = (r_1, r_2, \dots, r_N)'$ is used to denote the vector of random returns.

The optimal portfolio according to Markowitz is the one that minimizes the *variance* of the portfolio return,

$$\mathrm{Var}[R(\boldsymbol{\pi})] = \sum_{i,j=1}^{N} \pi_i \pi_j \left\langle (r_i - \mu_i)(r_j - \mu_j) \right\rangle = \sum_{i,j=1}^{N} \pi_i \pi_j \Sigma_{ij} \ , \tag{2}$$

subject to the constraint of a given expected portfolio return μ_P

$$\mu_P \equiv \langle R(\boldsymbol{\pi}) \rangle = \sum_{i=1}^{N} \pi_i \langle r_i \rangle = \sum_{i=1}^{N} \pi_i \mu_i \ . \tag{3}$$

In (2), $\Sigma = (\Sigma_{ij})$ is the covariance matrix of asset returns.

To put a scale to the problem, one usually imposes the normalization constraint

$$\boldsymbol{\pi}' \mathbf{1} \equiv \sum_{i=1}^{N} \pi_i = 1 \ . \tag{4}$$

Here $\mathbf{1} = (1, 1, \dots, 1)'$ denotes the N dimensional vector with all components equal to 1. The minimization problem is solved using the method of Lagrange multipliers to take the constrains of expected return and normalization into account, i.e. one looks the stationary point of the Lagrangian

$$\mathcal{L} = \frac{1}{2} \boldsymbol{\pi}' \Sigma \boldsymbol{\pi} - \lambda_1 (\boldsymbol{\pi}' \mathbf{1} - 1) - \lambda_2 (\boldsymbol{\pi}' \boldsymbol{\mu} - \mu_P) \tag{5}$$

w.r.t variations of the π_i, λ_1 and λ_2. Elementary linear algebra then entails that the optimal portfolio $\boldsymbol{\pi}^*$ takes the form

$$\boldsymbol{\pi}^* = \lambda_1 \Sigma^{-1} \mathbf{1} + \lambda_2 \Sigma^{-1} \boldsymbol{\mu} \ , \tag{6}$$

with actual values of the Lagrange parameters λ_1 and λ_2 determined by the constraints.

B. Translation into the Time-Domain

The Markowitz portfolio optimization problem allows a fairly straightforward translation into the time-domain. To formulate it, assume that $X = (X_t)_{t\in\mathbb{Z}}$ is the price process for a single traded asset. Let π_t denote the trading position that an investor takes on this asset at time t. As in the above we shall use the convention that $\pi_t > 0$ represents a long position (buying the asset), whereas $\pi_t < 0$ represents a short position (selling the asset).

The return of a trading strategy $\boldsymbol{\pi} = (\pi_1, \pi_2, \dots, \pi_T)'$ over a finite time horizon of T time steps for a realization $\boldsymbol{x} = (x_1, x_2, \dots, x_T)'$ of the price process can be written as

$$R_T(\boldsymbol{\pi}|x_0) = \sum_{t=1}^{T} \pi_t(x_0 - x_t) \ . \tag{7}$$

In terms of these conventions the expected return μ_S of a trading strategy (conditioned on the initial price x_0) is

$$\mu_S = \langle R(\boldsymbol{\pi}|x_0)\rangle = \sum_{t=1}^{T} \pi_t(x_0 - \mu_t) = x_0 - \boldsymbol{\pi}'\boldsymbol{\mu} \ , \tag{8}$$

where we have restricted ourselves in the second step to normalized trading strategies satisfying $\boldsymbol{\pi}\mathbf{1} = 1'$, and where $\mu_t = \langle x_t\rangle$ denotes the expected price at time t.

It is worth remaking at the outset that X could alternatively (and perhaps even more appropriately in the present context) be thought of as the log-price process, in which case $R_T(\boldsymbol{\pi}|x_0)$ would be the log-return of the strategy $\boldsymbol{\pi}$. For the sake of simplicity and definiteness we shall stick to the language of price processes and returns in what follows.

An optimal trading strategy in the spirit of Markowitz would then be a strategy which minimize the (conditional) variance

$$\mathrm{Var}[R_T(\boldsymbol{\pi}|x_0)] = \sum_{t,t'=1}^{T} \pi_t\pi_{t'}\, \langle (x_t - \mu_t)(x_{t'} - \mu_{t'})\rangle = \sum_{t,t'=1}^{T} \pi_t\pi_{t'}\Sigma_{tt'} \ , \tag{9}$$

subject to the constraints of normalization $\boldsymbol{\pi}'\mathbf{1} = 1$ and given mean return $\boldsymbol{\pi}'\boldsymbol{\mu} = x_0 - \mu_S$. In (9), the matrix $\Sigma = (\Sigma_{tt'})$ now denotes the *auto*-covariance matrix of the price process.

The algebraic side of the problem of finding an optimal trading strategy is now formally fully equivalent to that of finding an optimal portfolio, and the optimal strategy $\boldsymbol{\pi}^*$ takes the form

$$\boldsymbol{\pi}^* = \lambda_1\Sigma^{-1}\mathbf{1} + \lambda_2\Sigma^{-1}\boldsymbol{\mu} \ , \tag{10}$$

with Σ now the *auto*-covariance matrix of the price process rather than the covariance matrix of portfolio returns. Actual values of the Lagrange parameters λ_1 and λ_2 are determined by the constraints as before.

It is well known, and indeed easily verified that the globally optimal solution which does not impose a restriction concerning the mean return is compactly given by

$$\boldsymbol{\pi}^*_{\mathrm{GO}} = \frac{\Sigma^{-1}\mathbf{1}}{\mathbf{1}'\Sigma^{-1}\mathbf{1}} \ . \tag{11}$$

The main problem facing, both, portfolio optimization à la Markowitz, and the mean-variance approach to finding optimal trading strategies is in the fact that covariance matrices of portfolio returns or auto-covariance matrices of price processes of traded assets are not known, but need to be *estimated* from empirical market data. The effects of sampling noise in such estimation processes are meanwhile well studied in the case of portfolio optimization. As mentioned in the introduction, various strategies to mitigate against such effects — typically guided by random matrix theory — have been investigated in the past.

By contrast, the corresponding random matrix theory for sample auto-covariance matrices that might be invoked for similar purposes for the problem of mean-variance formulations of optimal trading strategies has only recently become available [29]. We shall address the issue of sampling noise in empirical data and the use of spectral theory for the purpose of guiding the choice of "cleaning"-strategies for auto-covariance matrices of market data below in Sect. IV. Before that we investigate the effects of sampling noise for some synthetic processes where comparison with known true auto-covariance matrices is possible.

III. RESULTS FOR SYNTHETIC PRICE PROCESSES

In this section we evaluate the theory developed in the previous section for synthetic price processes. We begin by taking these processes to be either white noise processes or auto-regressive processes of order 1, and then move on to look at the situation where price-*increments* are modelled as white-noise and auto-regressive processes, respectively. For the white noise and auto-regressive price processes, the true auto-covariance matrices are known, and analytical expressions for optimal trading strategies can be given. We then look at the effects of sampling noise, using *estimates* of auto-covariance matrices for various values of the ratio of $\alpha = T/M$ of the length T of the risk horizon (and thus the matrix dimension) and the sample size M used to determine these estimates. The analytical expressions for the true auto-covariance matrices correspond to the $\alpha \to 0$-limit in these results.

A. Synthetic Stationary Price Processes

We first consider a price process with fluctuations around the trend, $\delta x_t = x_t - \mu_t$ taken to be a Gaussian white noise process, i.e. $\delta X_t \sim \mathcal{N}(0, \sigma^2)$. The true auto-covariance matrix in this case is proportional to the unit matrix, i.e. $\Sigma_{t,t'} = \sigma^2 \delta_{t,t'}$.

The globally optimal strategy (11) for a time horizon of length T in this case is then readily found to be

$$\pi^*_{t,\mathrm{GO}} = \frac{\sigma^{-2}}{\sum_{t=1}^{T} \sigma^{-2}} = \frac{1}{T} \ . \tag{12}$$

Thus, for a white noise process with variance σ^2 the optimal strategy $\boldsymbol{\pi}^*_{\mathrm{GO}} = (1/T, 1/T, ..., 1/T)'$ is uniform over the time horizon T, and independent of the variance of the price process. The analogous result for a Markowitz portfolio of uncorrelated assets is, of course, well known.

Let us next assume that price fluctuations around the trend are described by an AR(1) process, i.e. an auto-regressive process of order 1 of the form

$$\delta X_t = a\, \delta X_{t-1} + \left(\sqrt{1-a^2}\right) \xi_t \ , \tag{13}$$

in which $\xi_t \sim \mathcal{N}(0, 1)$; for simplicity, we have normalized the process to exhibit fluctuations of variance 1. The parameter a in (13) is required to satisfy $|a| < 1$ for fluctuations to be stationary. The auto-covariance function of this process is known to be given by

$$\gamma(i) = \mathrm{Cov}[\delta X_t \delta X_{t-i}] = a^{|i|} \ . \tag{14}$$

The auto-covariance matrix evaluated for a finite time horizon of length T is thus a Toeplitz matrix of the form

$$\boldsymbol{\Sigma} = \begin{pmatrix} 1 & a & a^2 & & \cdots & a^{T-1} \\ a & 1 & a & a^2 & & \vdots \\ a^2 & a & 1 & a & \ddots & \\ & a^2 & a & 1 & \ddots & a^2 \\ \vdots & & \ddots & \ddots & \ddots & a \\ a^{T-1} & \cdots & & a^2 & a & 1 \end{pmatrix} \ . \tag{15}$$

It's inverse is a tri-diagonal matrix given by

$$\boldsymbol{\Sigma^{-1}} = \frac{1}{1-a^2} \begin{pmatrix} 1 & -a & 0 & \cdots & \cdots & 0 \\ -a & 1+a^2 & -a & \ddots & & \vdots \\ 0 & \ddots & \ddots & \ddots & \ddots & \vdots \\ \vdots & \ddots & \ddots & \ddots & \ddots & 0 \\ \vdots & & \ddots & -a & 1+a^2 & -a \\ 0 & \cdots & \cdots & 0 & -a & 1 \end{pmatrix} \ . \tag{16}$$

The globally optimal strategy (11) for a time horizon of length T in this case is then given by

$$\boldsymbol{\pi}^*_{\mathrm{GO}} = \lambda_1(1, 1-a, \ldots, 1-a, 1)' \ , \tag{17}$$

with $\lambda_1 = [2+(T-2)(1-a)]^{-1}$ fixed by the normalization-constraint $\boldsymbol{\pi}'\mathbf{1} = 1$. In this case the globally optimal trading strategy turns out to be uniform apart from the two boundary terms. The white noise result is clearly recovered as the $a \to 0$-limit of the present result for the AR(1) process as it should.

Solutions with constraints on the expected return can be given in closed form as well; they are simply obtained by inserting (16) into (10), with Lagrange parameters obtained by solving a pair of linear constraint-equations; details will of course depend on assumptions concerning the drift, and we refrain from writing them down explicitly.

Fig. 1 shows optimal strategies for an AR(1) price process with parameter $a = 0.8$, both for the global optimum as well as for cases with non-zero mean returns imposed. As can be seen from the figure, increasing the expected strategy return from $\mu_S = 4.0 \times 10^{-4}$ to $\mu_S = 1.0 \times 10^{-3}$ changes the optimal strategy (10) from one that is monotone decreasing over the risk-horizon to one which is monotone *increasing*, and starting in fact with a (short-)selling position at the initial time-step $t = 1$.

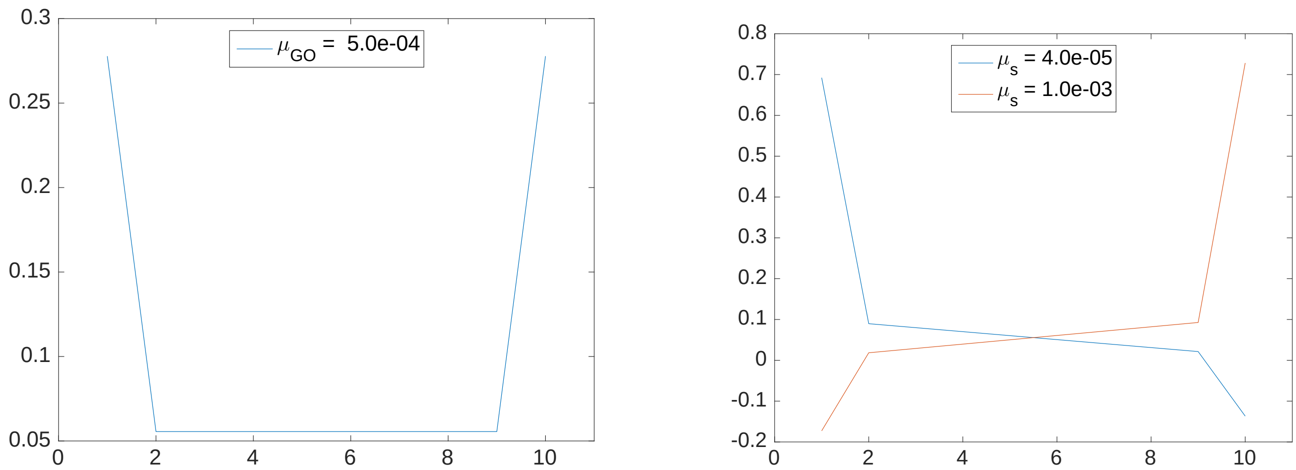

FIG. 1: Left panel: Globally optimal trading strategy for an AR(1) price process with $a = 0.8$ over a risk horizon of $T = 10$ time steps. Right panel: optimal strategies for a process with the same parameter a and a linear drift of the form $\mu_t = 10^{-4}t$, imposing expected strategy returns of $\mu_S = 4 \times 10^{-5}$ (blue solid line) and $\mu_S = 1 \times 10^{-3}$ (solid orange line).

B. Synthetic Price Processes with Stationary Increments

The stationarity assumption for the price process used in the previous subsection is clearly unrealistic, and there is obviously need to go beyond that, if the methods discussed in the present investigation are to stand the chance to be in anyway useful in practice.

However, once the realm of stationarity is left, some structure is needed on a different level in order to make operational sense of estimating auto-covariance functions and the corresponding auto-covariance matrices defined over a finite time horizon. The structure we shall rely on here is based on the assumption that (fluctuations of) price-process can be described as having *stationary increments*. If one adopts the reading that the processes considered here are actually log-price processes, the assumption of stationarity of their increments is actually a popular assumption in much of Mathematical Finance.

In what follows we assume that the (log-) price process $X = (X_t)$ exhibits stationary increments, i.e. that

$$X_t = X_{t-1} + Y_t \tag{18}$$

with $Y_t = \langle Y_t \rangle + \delta Y_t = \mu_t - \mu_{t-1} + \delta Y_t$ with zero-mean fluctuations δY_t. In terms of these conventions we can write the return of a strategy $\boldsymbol{\pi} = (\pi_t)$ for a given realization $\boldsymbol{x}$ as

$$R_T(\boldsymbol{\pi}) = \sum_{t=1}^{T} \pi_t (x_0 - x_t) = \sum_{t=1}^{T} \pi_t \Big[(\mu_0 - \mu_t) - \sum_{\tau=1}^{t} \delta y_\tau \Big] \ . \tag{19}$$

The expected return is given by the first contribution on the r.h.s, while the variance is

$$\mathrm{Var}[R_T(\boldsymbol{\pi})] = \sum_{t,t'=1}^{T} \pi_t \pi_{t'} \Big[\sum_{\tau=1}^{t} \sum_{\tau'=1}^{t'} \langle \delta y_\tau \delta y_{\tau'} \rangle \Big] \ . \tag{20}$$

This is of the same structure as (9), with the auto-covariance matrix $\Sigma \equiv \Sigma^X = (\Sigma^X_{t,t'})$ of the non-stationary price process expressed in terms of the auto-covariance matrix $\Sigma^Y = (\Sigma^Y_{t,t'})$ of the process of price increments as

$$\Sigma^X_{t,t'} = \sum_{\tau=1}^{t} \sum_{\tau'=1}^{t'} \langle \delta y_\tau \delta y_{\tau'} \rangle = \sum_{\tau=1}^{t} \sum_{\tau'=1}^{t'} \Sigma^Y_{\tau,\tau'} \ . \tag{21}$$

This relation between the auto-covariance matrices of process and the corresponding process of increments can be compactly expressed in matrix form as

$$\Sigma^X = P \Sigma^Y P' \ , \tag{22}$$

where P is a lower triangular constant matrix of ones,

$$P = \begin{pmatrix} 1 & 0 & 0 & \dots & 0 \\ 1 & 1 & 0 & \dots & 0 \\ 1 & 1 & 1 & \dots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 1 & 1 & 1 & \dots & 1 \end{pmatrix} \ . \tag{23}$$

The mean variance approach to strategy optimization then yields optimal trading strategies of the form (10), with the auto-covariance matrix $\Sigma = \Sigma^X$ of the price process expressed in terms of the auto-covariance matrix Σ^Y of the process of stationary increments according to Eq. (22)

Taking the process of price increments to be a white noise process $\delta Y_t \sim \mathcal{N}(0, \sigma^2)$, we have have $\Sigma^Y_{t,t'} = \sigma^2 \delta_{t,t'}$ so $\Sigma^{-1} = \sigma^{-2} (PP')^{-1}$, where $(PP')^{-1}$ is found to be of tridiagonal form,

$$(PP')^{-1} = \begin{pmatrix} 2 & -1 & 0 & 0 & \dots & 0 \\ -1 & 2 & -1 & 0 & \dots & 0 \\ 0 & -1 & 2 & -1 & \dots & 0 \\ \vdots & & \ddots & \ddots & \ddots & \vdots \\ 0 & 0 & \dots & -1 & 2 & -1 \\ 0 & 0 & \dots & & -1 & 1 \end{pmatrix} \ . \tag{24}$$

The globally optimal strategy (11) in this case is then simply

$$\boldsymbol{\pi}^*_{\mathrm{GO}} = (1, 0, 0, \dots 0)' \ , \tag{25}$$

i.e., it consists of taking a single long position at the initial time step.

If we assume an AR(1) process of the form (13) of a white noise process for the fluctuations of the process of price increments, i.e.

$$\delta Y_t = a \, \delta Y_{t-1} + \Big(\sqrt{1 - a^2} \Big) \xi_t \ , \tag{26}$$

then it is Σ^Y which is given by Eq. (15); it turns out that $\Sigma^{-1} = (P \Sigma^Y P')^{-1}$, too, can be evaluated in closed form, giving

$$\mathbf{\Sigma^{-1}} = \frac{1}{1-a^2}\begin{pmatrix} C & -A^2 & a & 0 & \cdots & \cdots & & \cdots & 0 \\ -A^2 & 2B & -A^2 & a & 0 & & & & \vdots \\ a & -A^2 & 2B & -A^2 & a & 0 & & & \vdots \\ 0 & \ddots & \ddots & \ddots & \ddots & \ddots & \ddots & & \vdots \\ \vdots & \ddots & & & & & & \ddots & \vdots \\ \vdots & & \ddots & \ddots & \ddots & \ddots & \ddots & \ddots & 0 \\ \vdots & & & 0 & a & -A^2 & 2B & -A^2 & a \\ \vdots & & & & 0 & a & -A^2 & C & -A \\ 0 & \cdots & \cdots & \cdots & \cdots & 0 & a & -A & 1 \end{pmatrix} .$$

in which we use the abbreviations $A = 1 + a$, $B = 1 + aA$ and $C = 1 + A^2$.

In this case the globally optimal strategy (11) is of the form

$$\boldsymbol{\pi}^*_{\mathrm{GO}} = (1 + a, -a, 0, \ldots, 0)' , \tag{27}$$

i.e. it consists of taking a single long position at the first time-step, which is then partially offset by a short position at the second time step if $a > 0$, whereas it is followed by a further long position if successive price increments are *anti*-correlated ($a < 0$). Note that the solution for white noise increments is correctly recovered as the $a \to 0$-limit of the AR(1) results.

Once more, solutions with constraints on expected returns can be given in closed form; in analogy to the procedure described for the case of stationary price processes, they are obtained by inserting (16) into (10), with Lagrange parameters obtained by solving a pair of linear constraint-equations.

We find, and shall demonstrate below that the procedure predicts non-trivial changes of strategy, as constraints on expected returns are varied. Once more, details will depend on assumptions concerning the drift, and we refrain from producing explicit equations here. We will report our analytical results alongside numerical results which take sampling errors arising from finite sample fluctuations on *estimated* auto-covariance matrices into account

C. The Effects of Sampling Noise

Having analytical results for synthetic price processes available allows one to estimate the effects of sampling noise on optimal strategies and on risk return profiles. In practice, the analytic structure of an underlying price process will not be known, and auto-covariance matrices will have to be estimated on the basis of finite samples, i.e. the design of optimal strategies will have to be based on *sample auto-covariance matrices* $\hat{\Sigma}$.

For a stationary price process, samples taken along a realization of the process can be taken to define the elements of $\hat{\Sigma}$ via

$$\hat{\Sigma}_{t,t'} = \frac{1}{M-1}\sum_{\mu=1}^{M} \delta x_{t+\mu}\delta x_{t'+\mu} . \tag{28}$$

This procedure introduces sampling noise; estimated auto-covariance matrix elements $\hat{\Sigma}_{t,t'}$ will exhibit $\mathcal{O}(M^{-1/2})$ fluctuations about their corresponding true counterparts $\Sigma_{t,t'}$. When assessing the effects of sampling noise via the influence on spectra, one expects the relevant parameter to be the aspect ratio $\alpha = T/M$, i.e. the ratio of the number of time-lags considered and the sample-size used to estimate matrix elements. We shall use this parameter in what follows to parametrize the influence of sampling noise, with the $\alpha \to 0$-limit corresponding to the situation without sampling noise, i.e. with true asymptotic auto-covariances known.

If the price process is not stationary, but has stationary increments, one can use Eqs. (21) and (22) to express the auto-covariance matrix Σ^X of the price process in terms of the auto-covariance matrix Σ^Y of the process of

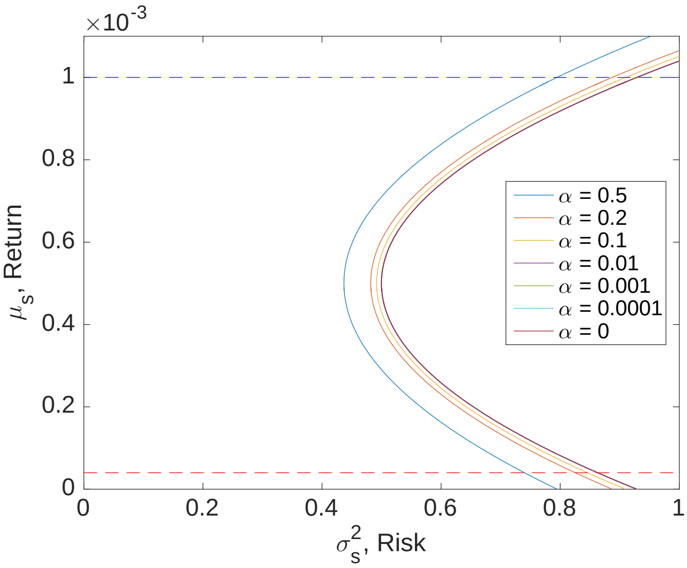

FIG. 2: Risk-return profile for an AR(1) price process with the same parameters as in Fig. 4 for various levels of sampling noise parameterized by α. Results are obtained by averaging over 10^7 samples as in Fig.3. Note in particular that sampling noise leads to an *under-estimation of risk*. The two horizontal dashed lines indicate two values of the target return for which optimal trading strategies are reported in Fig 4 below.

price increments. For the latter it is legitimate to use an estimator by sampling along a realization, so one can define $\hat{\Sigma}^X$ via

$$\hat{\Sigma}^Y_{t,t'} = \frac{1}{M-1}\sum_{\mu=1}^{M} \delta y_{t+\mu}\delta y_{t'+\mu} \tag{29}$$

and

$$\hat{\Sigma}^X = P\hat{\Sigma}^Y P' \ . \tag{30}$$

In Fig. 2 we show the risk-return profile for the case of an AR-1 price process for various aspect ratios α, ranging from $\alpha = 0.5$ down to $\alpha = 10^{-4}$, with the noise-free case $\alpha = 0$ also included. Note that sampling noise leads to a systematic underestimation of risk, though results quickly approach the noise-free limit as α becomes small.

Fig. 3 exhibits the weights of the globally optimal (minimum risk) trading strategy for this process, while Fig. 4 gives weights of optimal trading strategies for two different values of the target return (indicated by the two horizontal dashed lines in Fig. 2. In this case we assume a small underlying drift $\mu_t = 10^{-4}t$ of the underlying price process. It is noticeable that an increase in the required target return leads to a qualitative change of the optimal strategy, with the larger target return requiring to take an initial short position at the beginning of the trading period.

Turning to the situation where we use an auto-regressive process to describe the statistics of price *increments*, we see from a comparison of Figs. 5 and 2 that risk levels are significantly larger compared to the situation where the same underlying process describes the fluctuations of the price process itself.

This concludes our collection of results for synthetic price processes, where the underlying true auto-covariances are known. We now turn to applying the framework to empirical data, where this is not the case.

IV. EMPIRICAL DATA

In what follows we apply our framework to empirical data, using daily adjusted close data of the S&P500, spanning the period 03. Jan 1950 to 20. Apr 2015.

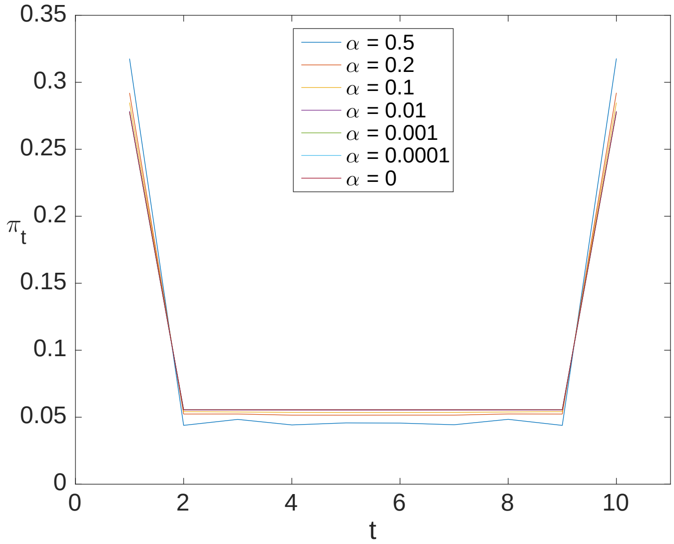

FIG. 3: Globally optimal trading strategies for an AR(1) price process with $a = 0.8$ over a risk horizon of $T = 10$ time steps, using *estimated* auto-covariance matrices. Data are shown for various values of the ratio $\alpha = T/M$ of risk horizon and sample size M used to estimate auto-covariances according to Eq. (28): optimal strategies (with solid lines as guides to the eye) are obtained by averaging over 10^7 samples. Standard deviations are not shown; they rapidly decrease with α. Results obtained for the *true* auto-covariance function (the $\alpha \to 0$-limit) are included for comparison. Note that average strategies obtained for finite samples are very close to the $\alpha = 0$ results.

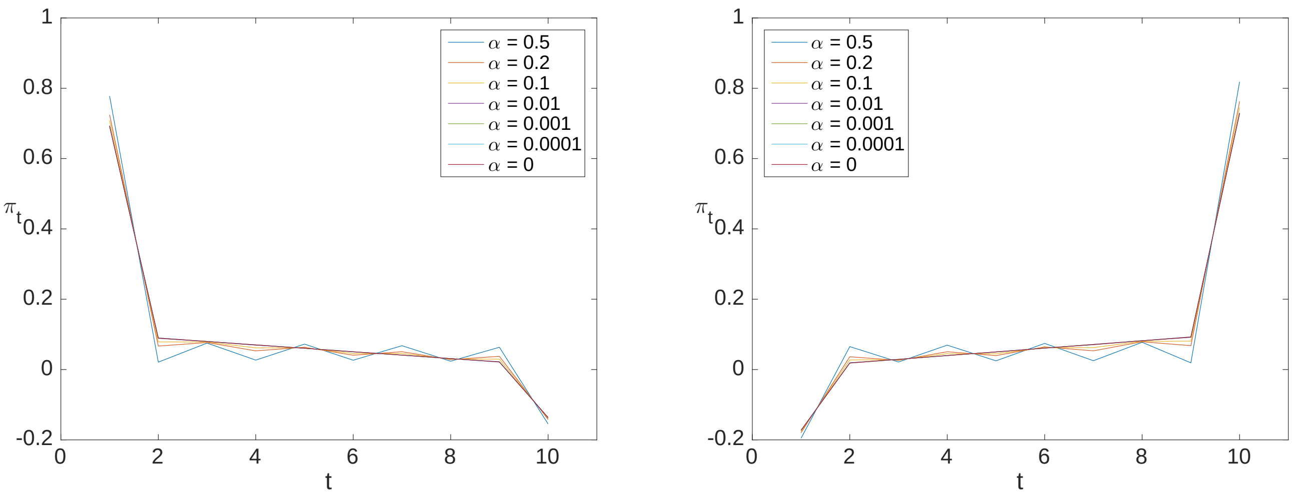

FIG. 4: Left panel: Optimal strategies for an AR(1) process with $a = 0.8$ and a linear drift of the form $\mu_t = 10^{-4}t$ as in Fig. 1, with imposed expected strategy return of $\mu_S = 4 \times 10^{-5}$. Shown are average trading strategies for various levels of sampling noise paramterized by non-zero α, obtained by averaging over 10^7 samples. Average results are close to those obtained using true asymptotic auto-covariance matrices in the $\alpha \to 0$-limit, which are included for comparison. Right panel: optimal trading strategies for an AR(1) process with the same parameters as in the left panel, but now with $\mu_S = 1 \times 10^{-3}$.

This is perhaps the point to notice that we are not advocating that using the variance of trading strategy returns constitutes the best way of capturing risk in real market data. Indeed, given that market returns are known to have fat-tailed distributions, variance can at best be regarded as a proxy for risk. Howevever, our primary goal here is not to explore a wider family of possible risk measures, but rather to define a reformulation of the popular mean-variance optimization strategy in the time domain, and to begin investigating its properties.

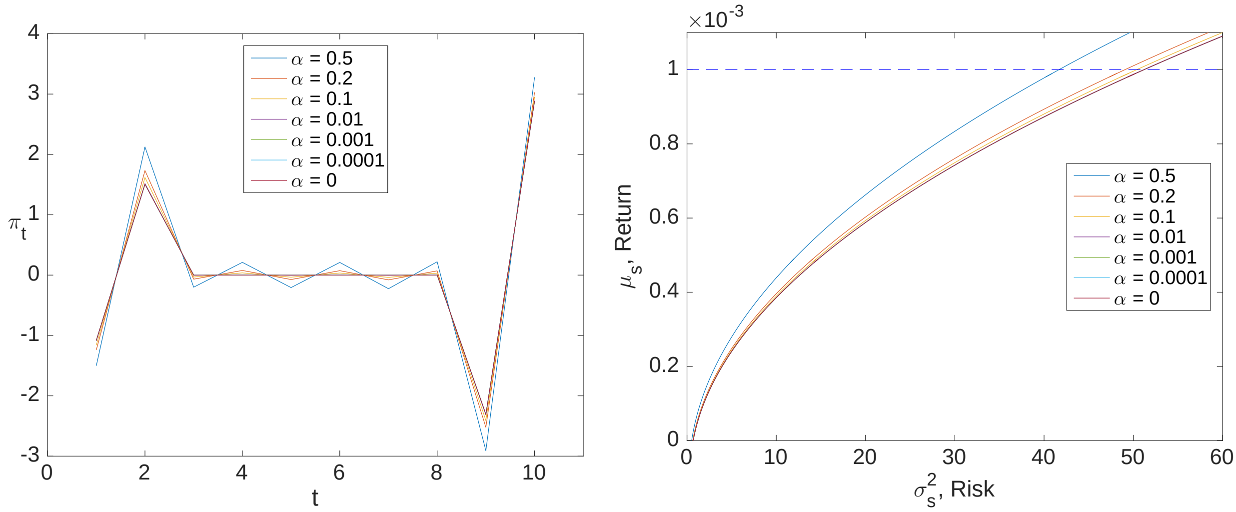

FIG. 5: Left: Optimal strategies for a setup where the fluctuations of the *price-increments* are described by an AR(1) process with $a = 0.8$; a linear linear drift of the form $\mu_t = 10^{-4}t$ is assumed for the price process, and an expected strategy return of $\mu_S = 1 \times 10^{-3}$ is imposed. Shown are average trading strategies (solid lines) obtained by averaging over 10^7 samples for various levels of sampling noise paramterized by non-zero α. Average results are close to those obtained using true asymptotic auto-covariance matrices in the $\alpha \to 0$-limit, which are included for comparison. Right: risk-return profile for this setup, with the horizontal dashed line indicating the expected strategy return imposed in the data of the left panel. The right panel should be compared with Fig 2, which exhibits the risk return profile for an AR-1 price process.

A. The Spectrum of the S&P500 Auto-Correlation Matrix

Before turning to the evaluation of optimal trading strategies and risk-return profiles we shall have a look at the spectrum of the auto-covariance matrices of the data, taking time windows of $T = 50$, and sample sizes of $M = 100$, hence $\alpha = 0.5$. Auto-covariance matrices of the price process are obtained as described in Sect. III C, by first evaluating auto-covariances of the return process, assuming stationarity across individual sample-windows. In order to obtain a meaningful statistics across the entire data set, we transform the return series in each time window to exhibit unit-variance increments, and then obtain auto-covariances of the thus normalized price process using the transformation Eq. (30).

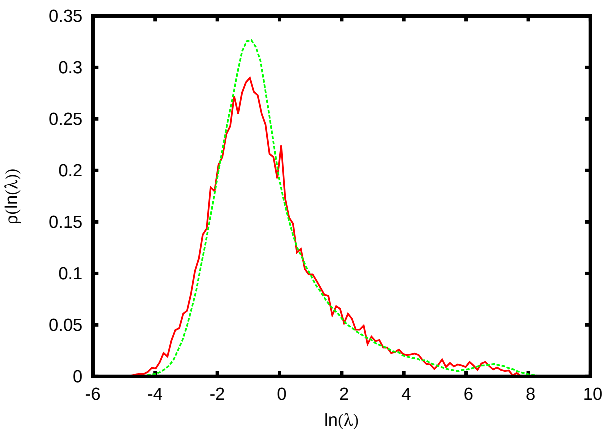

FIG. 6: Spectrum of the sample auto-covariance matrix of the S&P500, normalized as described in the main text, using $T = 50$ time lags and an aspect ratio $\alpha = 0.5$, i.e. samples of size $M = 100$ to define the auto-covariances (red full line). Also shown is a comparison with the spectrum of an auto-covariance matrix for a price process with *independent* unit variance increments (green dashed line). The two are remarkably close.

As can be seen in Fig. 6, where we plot the density of logarithms of eigenvalues, the spectrum is *very broad*, spanning several orders of magnitude. For comparison we include the spectrum for a process with *independent* unit variance increments using the same values of T and M, and we notice that the two are remarkably close. This is not completely unanticipated, as it is one of the widely reported 'stylized facts' in the field that return-series have very short correlation-times. We will use this type of spectral comparison below to inform the auto-covariance matrix cleaning strategy that we will use for the purpose of noise reduction.

B. Optimal Trading Strategies and Auto-Covariance Matrix Cleaning

In Fig 7 we report the risk-return characteristics for optimal trading strategies on the S&P500, using sample-auto-covariance matrices of $T = 50$ time lags, and sample size $M = 100$ as in Fig. 6. We report results obtained for auto-covariance matrices, as measured via Eqs. (29) and (30), and compare them with results obtained by applying a cleaning strategy to these, which we shall describe below. We use *realized* returns defined by linear trends in each data window to compute risk-return profiles, and use conventions for in-sample risk, true risk and and out-of-sample risk as e.g. in [31], taking the average auto-correlation matrix across the entire time series as a proxy for the true auto-correlation. Note that the reduction of risk that can be obtained through cleaning is substantial.

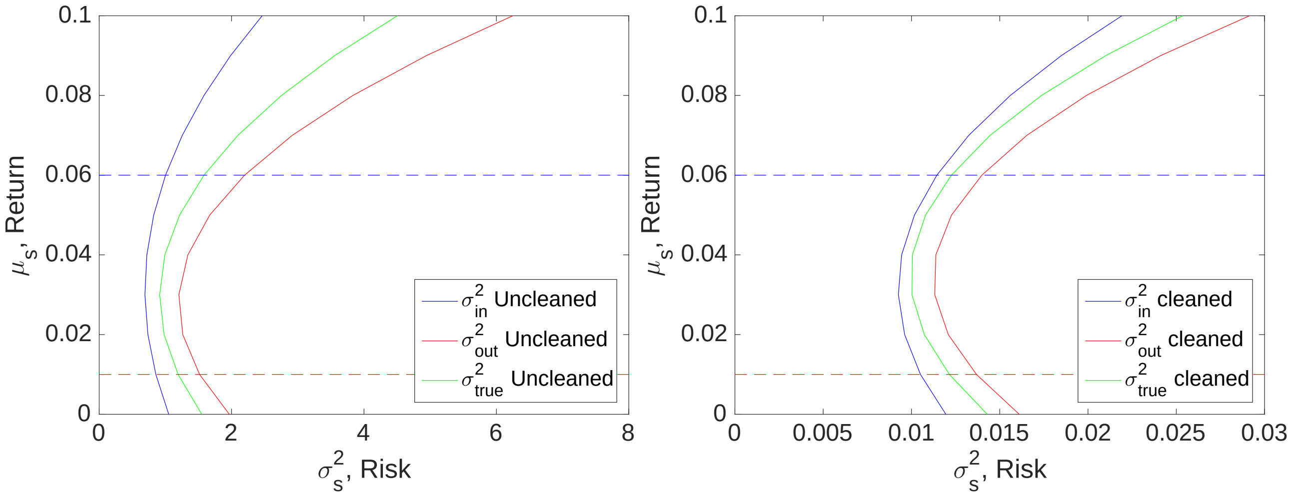

FIG. 7: Risk-return profile of optimal trading strategies on the S&P500 data. Left: risk-return profile obtained from measured auto-covariance matrices. Right: risk-return profile obtained using cleaned versions of auto-covariance matrices. Horizontal dashed lines denote target strategy returns μ_S for which optimal strategies are reported in Fig. 8 below.

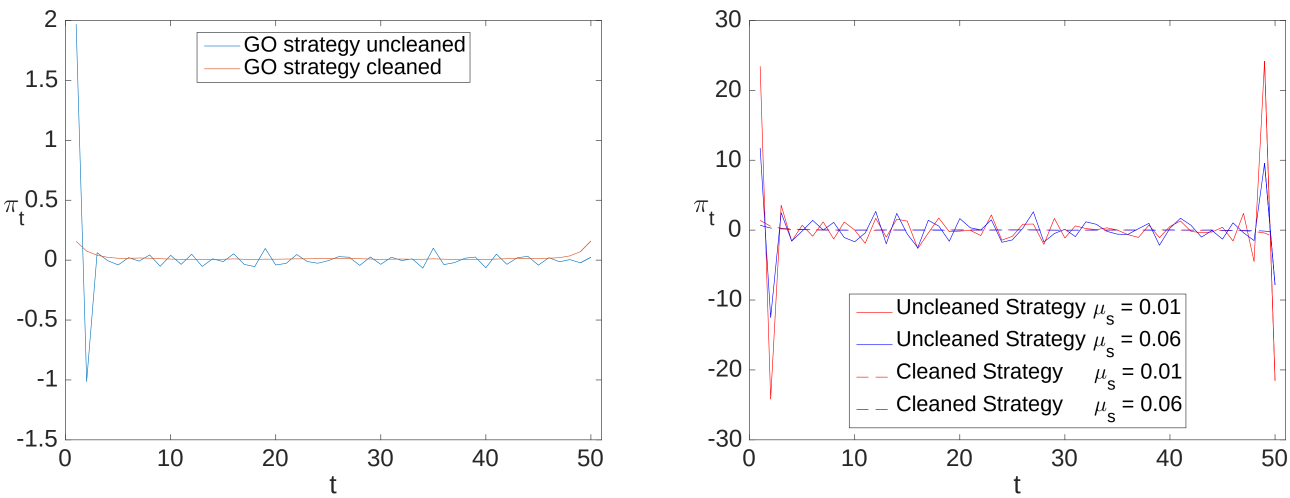

FIG. 8: Left: globally optimal strategy for the S&P500, showing both results before and after cleaning. Right: optimal strategies for the two target returns of $\mu_S = 0.01$ and 0.06 indicated in Fig. 7 above.

Fig. 8 exhibits optimal trading strategies for the S&P500, showing both the minimal risk solution and risk-optimal solutions for two different non-zero target strategy returns. Apart from the effect of reducing risk, we find that the effect of cleaning is also to create strategies that are "smoother" than those obtained without cleaning.

Let us finally turn to the cleaning strategy that is used to obtain the data described above. In the context of *covariance* matrices of financial data, strong similarities were observed between empirical correlation matrix spectra and the Marčenko-Pastur law expected for high-dimensional uncorrelated data. One of the cleaning strategies that has been suggested due to such similarities is referred to as 'clipping' [15, 31]. It analyses correlation matrices by performing a spectral decomposition, and regards the bulk of a sample correlation matrix spectrum, which resembles the Marčenko-Pastur law, as noise. It then transforms correlation matrices by keeping large eigenvalues outside the bulk, and replacing those in the bulk by their average, thereby avoiding small eigenvalues in the transformed matrix.

In the present case, the phenomenology is rather different; there are no eigenvalues of the (normalized) sample auto-covariance matrices which can be regarded as lying significantly outside the bulk of the spectrum predicted for uncorrelated increments. So there would be no clear guidance coming from random matrix theory that could form the basis of a clipping-type procedure.

We therefore decided to apply a 'shrinkage' procedure to our data. To the best of our knowledge this procedure was first proposed by Stein [32], and has recently found renewed interest in the Mathematical Statistics [18, 24] and Econophysics [33] communities.

Based on the observation reported in Fig 6 that the (normalized) auto-covariance spectra of the S&P500 and of a synthetic process with independent increments are indeed rather similar, we apply the shrinkage procedure to the sample auto-covariance matrixes of the S&P500 increments $\hat{\Sigma}^Y$, shrinking them towards a target matrix D given by the diagonal matrix of *variances* of the increments (which would indeed describe a process of independent increments), i.e. towards $D = \mathrm{diag}(\{\hat{\Sigma}_{t,t}\})$, using the substitution rule

$$\hat{\Sigma}^Y \leftarrow \delta D + (1-\delta)\hat{\Sigma}^Y \ , \tag{31}$$

and transforming the shrunk $\hat{\Sigma}^Y$ thus obtained to define the cleaned estimate of $\hat{\Sigma}^X$ using the transformation Eq. (22). The proper value for the parameter δ in this procedure is determined from the data as described in [18, 24].

V. SUMMARY AND DISCUSSION

To summarize, in the present paper we have looked at a reformulation of Markowitz' mean-variance optimization in the time domain to obtain optimal trading strategies for a single traded asset over a finite discrete time horizon. Using simple linear algebra, one obtains such optimal trading strategies as sequences of buy, hold, and sell instructions for that asset, which minimize the market fluctuations of the return generated by this sequence of instructions over a given time horizon, subject to suitable constraints. The procedure requires the auto-covariance matrix of the price process (and estimates for expected prices) during the risk horizon as input.

We investigated this problem for a number of synthetic price processes, taken to be either second order stationary or be described by second order stationary increments. Analytic expressions are given for the cases where the price and the return processes are described by i.i.d. or by auto-regressive fluctuations.

We compare analytic solutions with numerical results for situations where auto-covariance matrices have to be estimated from finite samples, which is the situation typically encountered in practice. For the synthetic processes for which true auto-covariance matrices are known the effects of sampling noise on optimal strategies and on risk-return profiles can thus be quantitatively assessed. We find that in general sampling noise leads to an underestimation of risk, but that asymptotic results are well approximated when samples used to estimate auto-covariance matrices are sufficiently large. A ratio $\alpha = T/M < 0.1$, i.e. sample sizes then times the length of the risk-horizon appears to be desirable from this point of view.

From the financial point of view on the other hand, it is always desirable to use time series as short as possible for estimation, to avoid letting (possibly) outdated data influence current trading strategies. Small samples, however, increase the effects of sampling noise, and it is for this reason that cleaning strategies have an important role to play. Looking at the S&P500 data, we found that (normalized) auto-covariance spectra closely resemble those one would expect for price processes with independent increments, and it is this observation that motivates our choice of target matrix within a shrinkage cleaning strategy.

We observe that auto-covariance matrix cleaning gives rise to smoother trading strategies, and that it also leads to a reduction of risk in risk-return profiles.

A natural generalization of the present work would deal with a multi-period multi-asset version of a mean-variance formulation of optimal trading strategies. While some work has been done in this direction in the past (see. e.g. [30] and references therein) the solution presented in [30] remains somewhat formal, and restricted to the case without correlations in time. We are not aware of an investigation of the effects of sampling noise in the multi-period multi-asset case. Indeed the spectral theory for that case which would be useful to motivate and design cleaning strategies has not been developed as of now.

Another direction that could be pursued is to include higher moments of strategy-return distributions in measures of risk, in order to better capture risk in the presence of fat-tailed return distributions. The translation into the time-domain, as advocated in the present paper would in general involve k-point correlations of returns in time (where $k \geq 3$). Assessing sampling noise in such a situation would then clearly transcend the realm of random matrix theory

[1] E. Aurell and P. Muratore-Ginanneschi. Growth-optimal strategies with quadratic friction over finite-time investment horizons. *International Journal of Theoretical and Applied Finance*, 7:645–657, 2004.
[2] K. Muthuraman and S. Kumar. Multidimensional portfolio optimization with portional transaction costs. *Mathematical Finance*, 16:301–335, 2006.
[3] K. Muthuraman and H. Zha. Simulation-based portfolio optimization for large portfolios with transaction costs. *Mathematical Finance*, 18:115–134, 2008.
[4] A. W. Lynch and S. Tan. Multiple risky assets, transaction costs, and return predictability: Allocation rules and implications for us investors. *Journal of Financial and Quantitative Analysis*, 45:1015–1053, 2010.
[5] S. Basak and G. Chabakauri. Dynamic mean-variance asset allocation. *Review of Financial Studies*, 23:2970–3016, 2010.
[6] D.B. Brown and J.E. Smith. Dynamic portfolio optimization with transaction costs: Heuristics and dual bounds. *Management Science*, 57:1752–1770, 2011.
[7] N.B. Garleanu and L.H. Pedersen. Dynamic trading with predictable returns and transaction costs. *Journal of Finance*, 68:2309–2340, 2013.
[8] V. DeMiguel, X. Mei, and F.J. Nogales. Multiperiod portfolio optimization with many risky assets and general transaction costs. Available at SSRN: http://ssrn.com/abstract=2295345.
[9] H.M. Markowitz. Portfolio selection. *The Journal of Finance*, 7:77–91, 1952.
[10] G. Connor, L.R. Goldberg, and R.A. Korajczyk. *Portfolio Risk Analysis.* Princeton University Press, Princeton and Oxford, 2010.
[11] E. J. Elton, M. J. Gruber, S J. Brown, and W. N. Goetzmann. *Modern Portfolio Theory and Investment Analysis.* Wiley, New York, 2010.
[12] K. Back. *Asset Pricing and Portfolio Choice Theory.* OUP, 2010.
[13] J.P. Bouchaud and M. Potters. *Theory of Financial Risk and Derivative Pricing: From Statistical Physics to Risk Management.* Cambridge University Press, Cambridge, 2006.
[14] A. Meucci. *Risk and Asset Allocation.* Springer, 2005.
[15] L. Laloux, P. Cizeau, J.P. Bouchaud, and M. Potters. Noise dressing of financial correlation matrices. *Phys. Rev. Lett.*, 83:1467–1470, 1999.
[16] V. Plerou, P. Gopikrishnan, B. Rosenow, L.A.N. Amaral, and H.E. Stanley. Universal and non-universal properties of cross-correlations in financial time series. *Phys. Rev. Lett.*, 83:1471–1474, 1999.
[17] S. Pafka and I. Kondor. Noisy covariance matrices and portfolio optimization ii. *Physica A: Statistical Mechanics and its Applications*, 319:487–494, March 2003.
[18] O. Ledoit and M. Wolf. Honey, I shrunk the sample covariance matrix. *J. Portfolio Management*, 31:110–119, 2004.
[19] G. Papp, S. Pafka, M. A. Nowak, and I. Kondor. Random matrix filtering in portfolio optimization. *Acta Physica Polonica B*, 36(9):2757–2765, 2005.
[20] L. Laloux. M. Potters, J. P. Bouchaud. Financial application of random matrix theory: Old laces and new pieces. *Acta Phys. Pol. B*, 36:2767, 2005.
[21] V. Golosnoy and Y. Okhrin. Multivariate shrinkage for optimal portfolio weights. *The European Journal of Finance,*, 13:441–45, 2007.
[22] Y. Chen, A. Wiesel, and A.O. Hero. Shrinkage estimation of high dimensional covariance matrices. In *IEEE Intl Conf. on Acoust., Speech, and Signal Processing*, 2009.
[23] S. Still and I. Kondor. Regularizing portfolio optimization. *New Journal of Physics*, 12:075034, 2010.
[24] O. Ledoit and M. Wolf. Nonlinear shrinkage estimation of large-dimensional covariance matrices. *The Annals of Statistics*, 40:1024–1060, 2012.
[25] F. Caccioli, S. Still, M. Marsili, and I. Kondor. Optimal liquidation strategies regularize portfolio selection. *The European Journal of Finance*, 19:554–571, 2013.

[26] F. Caccioli, I. Kondor, and M. Marsili amd S. Still. l_p regularized portfolio optimization. http://arxiv.org/pdf/1404.4040.pdf, April 2014.
[27] J. Wishart. Generalized Product Moment Distribution in Samples. *Biometrika*, 20 A:32–52, 1928.
[28] V. A. Marčenko and L. A. Pastur. Distribution of eigenvalues for some sets of random matrices. *Math. USSR-Sb.*, 1:457–483, 1967.
[29] R. Kühn and P. Sollich. Spectra of empirical auto-covariance matrices. *EPL (Europhysics Letters)*, 99:20008, 2012.
[30] D. Li an and W.L. Ng. Optimal dynamic portfolio selection: Multiperiod mean-variance formulation. *Mathematical Finance*, 10:387–406, 2000.
[31] J.P. Bouchaud and M. Potters. Financial applications of random matrix theory: a short review. In G. Akemann, J. Baik, and P. Di Francesco, editors, *The Oxford Handbook of Random Matrix Theory*. Oxford University Press, Oxford, 2011.
[32] C. Stein. Inadmissibility of the Usual Estimator for the Mean of a Multivariate Distribution. *Math. Rev.*, 1:197–206, 1956.
[33] M. Tumminello, F. Lillo, and R. N. Mantegna. Shrinkage and Spectral Filtering of Correlation Matrices: a Comparison Via the Kullback-Leibler Distance. *Act. Phys. Pol.*, B 38:4079–4088, 2007.