

Liquidity risk and instabilities in portfolio optimisation

Fabio Caccioli,

*Department of Computer Science, UCL, London, WC1E 6BT, UK
and Systemic Risk Centre,
London School of Economics and Political Sciences, London, UK*

Imre Kondor,

Parmenides Foundation, Kirchplatz 1, 82049 Pullach, Germany

Matteo Marsili,

*Abdus Salam International Centre for Theoretical Physics,
Strada Costiera 11, 34151 Trieste, Italy*

Susanne Still,

*Information and Computer Sciences, University of Hawai'i at
Mānoa, 1680 East-West Road, Honolulu 96822, Hawai'i, USA*

January 7, 2016

Abstract

We show that including a term which accounts for finite liquidity in portfolio optimization naturally mitigates the instabilities that arise in the estimation of coherent risk measures on finite samples. This is because taking into account the impact of trading in the market is mathematically equivalent to introducing a regularisation on the risk measure. We show here that the impact function determines which regularizer is to be used. We also show that any regularizer based on the norm ℓ_p with $p > 1$ makes the sensitivity of coherent risk measures to estimation error disappear, while regularizers with $p < 1$ do not. The ℓ_1 norm represents a border case: its “soft” implementation does not remove the instability, but rather shifts its locus, whereas its “hard”

implementation (including hard limits or a ban on short selling) eliminates it. We demonstrate these effects on the important special case of Expected Short-fall (ES) that is on its way to becoming the next global regulatory market risk measure.

1 Introduction

Risk management of large investment portfolios presents challenges that go beyond the textbook example of Markowitz portfolio optimization. For instance, asset managers have to deal with liquidity risk [1] in addition to market risk: Liquidity is always finite in real markets. Even for highly liquid assets, such as large cap stocks, the volume of limit orders available in the order book at a given time is of the order 10^{-5} of market capitalization, corresponding to 1% of the daily traded volume [2]. This means that large orders have to be split and executed incrementally, with the consequence that execution prices become uncertain. Estimating liquidation prices requires indeed a good estimation of market-impact. Market-impact refers to the tendency of prices to respond to trading activity: Liquidation of large positions makes an impact on prices, and moves the market against the investor. The quantification of market-impact is currently a very active field of research. Empirical evidence [3, 4] strongly suggests that market-impact is a concave function of the order size, and theories have been proposed to explain such empirical evidence [3, 5]. The form of market-impact predicted by such theories was the basis of the proposal for impact-adjusted accounting put forward in [6], who consider the implication of market-impact for the risk of a levered investor and show that accounting for market-impact in risk management improves risk assessment.

Liquidity considerations should therefore be taken into account when optimizing portfolios. At first sight this problem seems independent of another difficulty of portfolio selection: that of the large estimation error. Our purpose here is to reveal a close link between these two problems.

Risk minimization for large institutional portfolios suffers from the curse of dimensionality: the number of different assets (the dimension of the measured return vector), N , is often comparable to, or even larger than the sample size (the length of the available time series), T . As such, portfolio optimization belongs to the realm of high dimensional statistics [7]. Empirical estimates of returns, covariances, etc. are unstable under sample fluctuations.

The instability has been identified by [8, 9] as an algorithmic phase transition [10], occurring at a critical value of the ratio N/T (depending on the risk measure in question). At this critical point the estimation error diverges with a universal

exponent, independent of the risk measure, the nature of the underlying fluctuations, or even whether these fluctuations are stationary or GARCH-like [11].

Expected Shortfall (ES) [12, 13] is on its way to becoming the next global regulatory market risk measure [14, 15]. The sensitivity of ES to estimation error was studied in [8, 9, 16], where the critical value of the ratio N/T was determined, at which the estimation error diverges and the risk function becomes unbounded from below. This instability is a common weakness [17] of all coherent risk measures [18].

Various methods have been proposed to address this problem [19–34], including a number of works that made use of one or another form of regularization [35–40]. Regularization, the standard way to deal with high dimensional problems, imposes penalties on large excursions of the estimate. In [41] we advocated the systematic use of regularization in portfolio selection to eliminate the instability. The combination of a given risk measure with a choice of regularizer must be adjusted to the nature of the problem. We illustrated the idea of regularized portfolio optimization (RPO) using the Expected Shortfall as risk measure and the ℓ_2 norm as regularizer [41]. The optimization problem thus obtained was shown [41, 42]¹ to be closely related to support vector regression [43].

The relationship between market impact and regularization is revealed by realizing that regularized portfolio optimization arises naturally when investors take into account the impact their trades make on the market [44]. We demonstrated (again using the specific example of ES) that the ℓ_2 norm regularizer corresponds to linear market impact, and that the resulting optimization problem is identical to RPO, eliminating the instability [44].

In this paper we extend the scope of this idea substantially by demonstrating that considering market impact in the investment strategy results in a regularized portfolio optimization *for all coherent risk measures*. We identify the various regularizers corresponding to the different impact functions. We furthermore prove that the instability of *all* coherent measures can be removed by regularizers based on the ℓ_p norm with $p > 1$, but not by regularizers with $p < 1$.

The ℓ_1 norm sits on the borderline: it may or may not remove the instability, depending on the sample, which means that on average it does not remove, but only shifts it. In order to prevent the instability, the ℓ_1 norm regularizer has to be implemented as a ban on short selling, or any hard constraint that imposes infinite penalty on solutions wandering outside a finite domain.

We illustrate the effect of various regularizers on the estimation error of Expected Shortfall, as it has acquired special significance recently since the Basel Committee

¹This relationship was first recognized in [42], but the budget constraint (which is an important element of the portfolio problem, but does not enter support vector regression) was omitted there.

proposed it as the next global regulatory market risk measure [14, 15]. The calculations can be carried through analytically in the limit of large Gaussian portfolios ($N, T \rightarrow \infty$, with their ratio fixed) by methods borrowed from the statistical physics of disordered systems.

The paper is organized as follows. In Section 2 we show how market impact considerations lead to regularization for all coherent risk measures, and display the correspondence between the various impact functions and regularizers. Section 3 analyzes the effect of various ℓ_p regularizers on estimation error, with particular focus on the ℓ_1 norm. We start by analyzing a simple toy example: the Maximal Loss risk function (the limiting case of ES) with only two assets and two data points. This serves to provide intuition. We move on to show that the ℓ_1 regularizer cannot remove the instability of any of the coherent risk measures, while an ℓ_p with $p > 1$ can. Section 4 illustrates the above ideas using the example of a large random portfolio optimized under historical ES, and demonstrates the shift of the instability as a result of ℓ_1 -regularization. The effect of ℓ_p , $p > 1$, regularizers is also discussed, with special emphasis on the case $p = 3/2$ corresponding to a square root-like market impact that is characteristic of the liquidation of large positions under normal market conditions. A combination of the ℓ_1 and ℓ_2 norms as a regularizer is also considered in this Section. The last Section is a short summary, followed by the Appendix, where the technical details of the calculations have been relegated to.

2 Risk measures, market impact and Regularized Portfolio Optimization

The problem we focus on is that of determining the optimal portfolio $\vec{w} = (w_1, \dots, w_N)$ of N assets that an investor should hold. We assume the investor has a fixed wealth W to invest and can buy assets today at current prices \vec{p}_{now} . Therefore, feasible portfolios satisfy the budget constraint

$$\vec{w} \cdot \vec{p}_{\text{now}} = \sum_{i=1}^N w_i p_{i,\text{now}} = W. \quad (1)$$

The optimal portfolio is the one that minimizes the risk of the portfolio \vec{w} at a future time, at prices \vec{p}_{est} . Here we assume that

$$\vec{p}_{\text{est}} = \vec{p}_{\text{now}} + \vec{X} - \eta \vec{\psi}(\vec{w}), \quad (2)$$

where \vec{X} is a random vector of returns, $\vec{\psi}(\vec{w})$ denotes the market impact function, and η is a proportionality factor.² The term $\eta\vec{\psi}(\vec{w})$ accounts for the finite liquidity of the market. One way to motivate it is to argue that, in order to derive a monetary value from the portfolio, it ought to be sold. Then this term accounts for effects of finite liquidity, which are currently of great interest [3, 5]. The liquidation of the portfolio \vec{w} will generally move prices “against” the trading activity (i.e. $\psi_i(\vec{w}) \geq 0$). Much recent work has been devoted to deriving quantitative estimates of the market impact function $\psi_i(\vec{w})$ (see e.g. [3, 5]). Most of this literature has been concerned with the single asset setting ($N = 1$), suggesting that $\psi_i(\vec{w}) \propto w_i^\gamma$ with γ taking values around 1/2.

The value of the portfolio can then be estimated by the cash flow that can be generated when the portfolio is liquidated:

$$V(\vec{w}) = \vec{w} \cdot \vec{p}_{\text{est}} \quad (3)$$

$$= \vec{w} \cdot \vec{p}_{\text{now}} + \vec{w} \cdot \vec{X} - \eta\vec{w} \cdot \vec{\psi}(\vec{w}). \quad (4)$$

The optimal portfolio should be chosen in such a way as to minimize the risk attached to this cash flow [44].

Risk measures will be discussed in section 2.1. For the moment we just remark that risk enters only in the term $\vec{w} \left(\vec{X} - \eta \cdot \vec{\psi}(\vec{w}) \right)$, because $\vec{w} \cdot \vec{p}_{\text{now}} = W$ is known. Hence, without loss of generality and for the sake of simplifying mathematical expressions, we will omit this term in future discussion. Also, for the sake of simplicity, we normalize current prices to $p_{i,\text{now}} = 1/N$, $\forall i$, so that the budget constraint becomes

$$\sum_{i=1} w_i = NW. \quad (5)$$

2.1 Coherent risk measures

A risk measure is a function of the random variable $Z = \vec{w} \cdot \vec{X}$, denoted by $\rho[Z]$ ³ representing a loss, i.e. it is a quantity that one would like to minimize. Here we focus on coherent risk measures, a broad class of quantitative measures for risk [18]. A risk measure is coherent if it satisfies the following axioms:

²Notation: we use uppercase letters for random variables. $\psi_i(\vec{w})$ denotes the i^{th} component of the vector $\vec{\psi}(\vec{w})$.

³The squared parenthesis denotes that this is an operator $\rho : \mathcal{X} \rightarrow \mathbb{R}$ that associates a real number with any random variable $Z \in \mathcal{X}$.

Normalization $\rho[0] = 0$.

Monotonicity if $Z_1 \leq Z_2$ a.s. then $\rho[Z_1] \geq \rho[Z_2]$

Sub-additivity $\rho[Z_1 + Z_2] \leq \rho[Z_1] + \rho[Z_2]$

Positive homogeneity if $a \geq 0$ then $\rho[aZ] = a\rho[Z]$

Translation invariance if a is a real constant then $\rho[Z + a] = \rho[Z] - a$

The failure of the widely used risk measure Value at Risk to satisfy subadditivity was the prime motivation for the introduction of these axioms [18]. The translation invariance property implies that adding a certain amount a (of cash) to a portfolio reduces risk by a .

2.2 Market impact and regularization

Because of the translation invariance property, we have for any coherent risk measure, evaluated on the *value*, V , of a portfolio (see Eq. 4)

$$\rho[V] = \rho[\vec{w} \cdot \vec{X} - \eta \vec{w} \cdot \vec{\psi}(\vec{w})] = \rho[\vec{w} \cdot \vec{X}] + \eta \vec{w} \cdot \vec{\psi}(\vec{w}). \quad (6)$$

The first term corresponds to the risk of the portfolio without market impact considerations. The second term, $\eta \vec{w} \cdot \vec{\psi}(\vec{w})$, is a monotonically increasing function of the weights, therefore it penalizes large positions, which means that *it acts as a regularizer*. It is very natural that anticipation of future liquidation should limit the sizes of one's positions.

In summary, if the investor minimizes the risk of the cash flow that could be generated by the liquidation of the portfolio, then considering the impact of the investor's own actions on the market naturally leads to Regularized Portfolio Optimization (RPO) *for all coherent risk measures*.

The correspondence between market impact functions and regularizers is as follows:

The bid-ask spread impact $\vec{\psi}(\vec{w}) = \text{sign}(\vec{w})$ leads to the use of the ℓ_1 -norm, which we can write as $\vec{w} \cdot \text{sign}(\vec{w}) = \|\vec{w}\|_1$. This applies to single trades, as well as to sequential trades executed in a very short time.

Square root impact Large positions are usually liquidated by chopping the execution into many small child orders that are executed sequentially. Such *meta-orders* typically have an impact obeying a square root law $\psi_i(\vec{w}) = \sqrt{|w_i|}$ [3]. This leads to a $p = 3/2$ regularizing term: $\eta \vec{w} \cdot \vec{\psi}(\vec{w}) = \eta \sum_i |w_i|^{3/2}$.

Cube root impact As an alternative to the square root market impact rule, the $|w_i|^{1/3}$ rule was derived by Kyle and Obizhaeva within the context of market microstructure invariance [45, 46]. This leads to a $p = 4/3$ regularizing term $\eta \vec{w} \cdot \vec{\psi}(\vec{w}) = \eta \sum_i |w_i|^{4/3}$.

Linear impact On longer time scales, when the sales are repeated over a long time, we expect the impact to be linear: $\vec{\psi}(\vec{w}) = \vec{w}$. This leads to the regularizing term $\eta \vec{w} \cdot \vec{w} = \eta \|w\|_2$, i.e. implies the use of the ℓ_2 norm.

In general $\psi_i(\vec{w}) = |w_i|^{p-1}$ is equivalent to a regularizing term $\eta \sum_i |w_i|^p$, and thus implies the use of the ℓ_p -norm in RPO:

$$\text{Market impact: } \psi_i(\vec{w}) = |w_i|^{p-1} \Leftrightarrow \text{Regularizer: } \ell_p(\vec{w})$$

3 Stability of ℓ_p regularized portfolio optimization

In this section we analyze the stability of ℓ_p norm based regularized portfolio optimization. We focus particularly on $p = 1$, as it plays a special role. In particular, we show that, for all coherent risk measures, if $p > 1$, then regularization based on the ℓ_p norm removes the divergence that plagues portfolio optimization in the undersampled regime, but this is not the case for $p \leq 1$.

Let us recall that the phenomenon of interest is the instability of the empirical estimate of the risk measure, when the sample size, T , is of the same order as the number of assets, N , in the portfolio, i.e. the dimensionality of the return vector. Assume that a sample, \mathcal{X} , consists of T observations, \vec{x}_t , $t = 1, \dots, T$, drawn i.i.d. from a true (but unknown) underlying probability distribution, p , and that there is a space of samples, \mathbf{X} , which come from the same underlying distribution. The empirical estimate of a risk measure can then be obtained by using the empirical distribution

$$\hat{p}(\vec{x}) = \sum_{t=1}^T \frac{\delta(\vec{x} - \vec{x}_t)}{T} . \quad (7)$$

The empirical risk, $\hat{\rho}$, computed on this distribution, depends on the specific sample $\mathcal{X} = \{\vec{x}_t\}_{t=1}^T$. For some samples, the function $\hat{\rho}[\vec{w}]$ is unbounded from below, which means that the portfolio optimization problem does not have a finite solution. This results in an instability of the risk measure.

To develop an intuition, we first consider the simple Maximal Loss (ML) risk measure, regularized by the ℓ_1 norm, for $N = 2$ stocks and $T = 2$ observations (Section 3.1).

We show in Section 3.2 that any coherent risk measure can become unstable to estimation error for ℓ_1 norm based regularization. The proof generalizes that given in [17] for unregularized portfolio optimization. The proof relies on the probability of the existence of dominating portfolios. Depending on the ratio N/T , this probability can be either very close to zero – making the instability a rare event that can be neglected – or very close to one. Sharp results for the threshold value of N/T that separates the two cases can be given in the limit $N, T \rightarrow \infty$ with fixed ratio, using methods from statistical mechanics (see Section 4).

3.1 ℓ_1 regularized ML – a toy example

To develop an intuition about the regularized portfolio optimization problem with an ℓ_1 regularizer, we first consider a simple risk measure called Maximal Loss (ML), defined as [47]

$$\text{ML}(\vec{w}) = \min_{w_i} \max_{t=1, \dots, T} \left[- \sum_i w_i x_{i,t} \right]. \quad (8)$$

Maximal Loss is the best combination of the worst losses. As such it is a minimax type risk measure, a special case of Expected Shortfall (ES). Expected Shortfall is the average loss beyond a high quantile. The threshold α beyond which the average loss ES is calculated has typical values in the range (0.95 to 1.00). Maximal Loss corresponds to $\alpha = 1 - 1/T$.

The regularized portfolio optimization problem then becomes

$$\min_{w_i} \left\{ \max_{t=1, \dots, T} \left[- \sum_i w_i x_{i,t} \right] + \eta \|\vec{w}\|_p \right\}, \quad (9)$$

where $\|\vec{w}\|_p$ is the ℓ_p norm. This can be re-written to incorporate the regularizer directly into the loss function, since the term $\|\vec{w}\|_p$ does not depend on the individual returns, measured at times $t = 1, \dots, T$. We can write the *Regularized Maximal Loss* as

$$\text{RML}(\vec{w}) = \min_{w_i} \max_{t=1, \dots, T} \left[- \sum_i w_i x_{i,t} + \eta \|\vec{w}\|_p \right] \quad (10)$$

Now, we investigate stability resulting from use of the ℓ_1 norm, $\|\vec{w}\|_1 = \sum_i |w_i|$. To that end, let us consider the simple case of two assets ($i = 1, 2$) and two data points taken at times $t = 1$ and $t = 2$. The loss associated with each time can be expressed as

$$l_t = -w x_{1t} - (1 - w) x_{2t} + \eta |w| + \eta |1 - w| \quad (11)$$

where we have taken the budget constraint into account by setting $w_1 = w$ and $w_2 = 1 - w$. Now we distinguish three different cases:

- if $w > 1$, then $l_t = w(x_{2t} - x_{1t} + 2\eta) - x_{2t} - \eta$
- if $0 < w < 1$, then $l_t = w(x_{2t} - x_{1t}) + \eta - x_{2t}$
- if $w < 0$, then $l_t = w(x_{2t} - x_{1t} - 2\eta) + \eta - x_{2t}$

The instability may arise in the following situations:

- if $w < 0$, $x_{21} - x_{11} > 2\eta$ and $x_{22} - x_{12} > 2\eta$ (the two straight lines, l_1 and l_2 , both have positive slope).
- if $w > 1$, $x_{21} - x_{11} < -2\eta$ and $x_{22} - x_{12} < -2\eta$ (the two straight lines, l_1 and l_2 , both have negative slope).

This means that in the absence of regularization the Maximal Loss is unbounded from below whenever one of the assets dominates the other (i.e. its return is higher at both times). Since both the loss function and the regularizer are piecewise linear functions, there is a threshold value η_c for any given realization of returns such that if $\eta > \eta_c$ the instability is removed (see Figure 1). However, the threshold is in general different for different instances of the problem, therefore if we consider the whole ensemble of possible instances, the regularizer with a fixed, finite value of η will only prevent the instability with a certain probability.

Let us calculate this probability by considering the ensemble of all possible realizations of returns. Assuming that $x_{i,t}$ are independent normal distributed random variables, the probability that the instability is not removed, is given by

$$P(\eta) = \left(\int_{2\eta}^{\infty} \frac{e^{-z^2/4}}{\sqrt{4\pi}} dz \right)^2 + \left(\int_{-\infty}^{-2\eta} \frac{e^{-z^2/4}}{\sqrt{4\pi}} dz \right)^2 = \frac{1}{2} \operatorname{erfc}^2(\eta). \quad (12)$$

Here the integration variable z is distributed as the difference of two independent Gaussian variables.

The quantity $P(\eta)$ is plotted in Figure 2.

This probability is positive for any finite η , which means that a "soft" implementation of the ℓ_1 regularizer cannot prevent the instability. The instability disappears only for $\eta \rightarrow \infty$, when the regularizer imposes an infinite penalty on solutions with the weights outside the interval $[0,1]$. This "hard" constraint is equivalent to a ban on short selling.

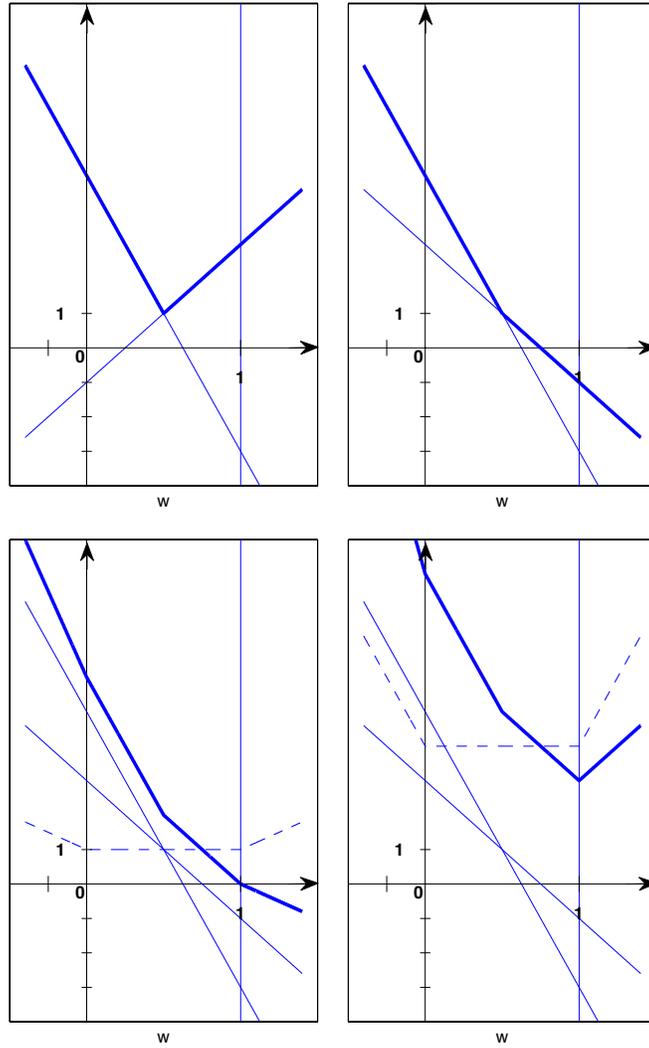


Figure 1: Maximal Loss with and without ℓ_1 regularizer for $N = T = 2$. Light solid lines represent losses at time $t = 1$ and $t = 2$. Bold lines represent the cost function of the respective optimization problem. The vertical solid line is drawn at $w = 1$. Dashed lines in the bottom panels represent the contribution of the regularizer to the cost function. **Top left.**: Maximal Loss with $N = T = 2$ when neither of the assets dominate, so a finite solution exists (the loss function has a finite minimum). **Top right.**: Asset 1 is dominating asset 2 here, the loss function does not have a finite minimum, and the solution runs off to infinity. **Bottom left.**: Same as top right, with an ℓ_1 regularizer added with a coefficient $\eta = 1$. This regularizer is too weak to prevent a run-away solution. **Bottom right.**: Same as top right, but now with a strong enough regularizer ($\eta = 4$) that keeps the solution finite. Note that the regularizer sets the weight of the dominating asset to one, and that of the dominated one to zero.

Whenever the original problem is unstable and the ℓ_1 regularizer removes the instability, it does so by setting the weight of the dominant asset to 1, and that of the dominated asset to 0. The well-known property of ℓ_1 – that it produces a sparse solution – means in the toy example considered here that it eliminates one of the assets.

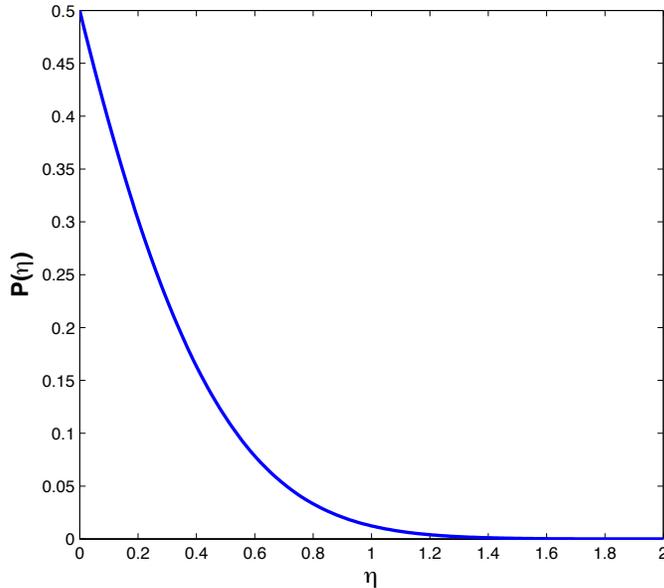


Figure 2: Probability that the ℓ_1 -regularized ML with two assets is unstable as a function of η .

The soft implementation of ℓ_1 fails to eliminate the instability, because the original loss function and the regularizer are piecewise linear functions. In contrast, a nonlinear regularizer, such as the ℓ_2 norm, removes the instability, as can be seen by considering the losses at times $t = 1, 2$:

$$l_t = -wx_{1t} - (1-w)x_{2t} + \frac{\eta}{2}w^2 + \frac{\eta}{2}(1-w)^2 \quad (13)$$

$$= \eta w^2 + (x_{2t} - x_{1t} - \eta)w - x_{2t} + \frac{\eta}{2} \quad (14)$$

The losses are now convex quadratic functions, therefore a finite and unique minimum for the Maximal Loss exists for all nonzero values of $\eta > 0$. The same holds true for any ℓ_p norm with $p > 1$, whereas a norm with $p < 1$ can not remove the instability.

3.2 Instability of ℓ_1 regularized coherent risk measures

We generalize the proof of the instability of coherent risk measures [17], which relies on the notion of dominant portfolios. The first step is to show that if such a dominant portfolio exists, then coherent risk measures are unbounded from below and the optimization problem is ill-defined. The second step is to show that the probability that dominant portfolios exist is bounded away from zero.

A vector \vec{u} is said to be a μ -dominant portfolio if *i*) $\vec{u} \cdot \vec{1} = 0$ and *ii*) $\vec{u} \cdot \vec{x}_t \geq \mu|\vec{u}|$ for all $t = 1, \dots, T$, where $|\vec{u}| = \sum_i |u_i|$ is the ℓ_1 norm.

Proposition: If a μ -dominant portfolio exists, then no finite ℓ_1 -regularized portfolio can exist for $\eta < \mu$.

Proof: The objective function that we wish to minimize is

$$\hat{\rho}[\vec{w} \cdot \vec{X}] + \eta|\vec{w}|. \quad (15)$$

For any normalized portfolio \vec{w} (i.e. $\vec{w} \cdot \vec{1} = 1$) consider the portfolio $\vec{w} + a\vec{u}$, where \vec{u} is a μ -dominant portfolio. This is a normalized portfolio (the sum of portfolio weights is one), because of *i*) in the definition of \vec{u} . Now

$$\begin{aligned} \hat{\rho}[(\vec{w} + a\vec{u}) \cdot \vec{X}] &\leq \hat{\rho}[\vec{w} \cdot X] + \hat{\rho}[a\vec{u} \cdot X] && \text{sub-additivity} \\ &= \hat{\rho}[\vec{w} \cdot X] + a\hat{\rho}[\vec{u} \cdot X] && \text{positive homogeneity} \\ &\leq \hat{\rho}[\vec{w} \cdot X] + a\hat{\rho}(\mu|\vec{u}|) && \text{monotonicity } (a > 0) \\ &= \hat{\rho}[\vec{w} \cdot X] - a\mu|\vec{u}| && \text{translation inv. + normalization.} \end{aligned} \quad (16)$$

To get to eq. (16), we used monotonicity with $Z_1 = \mu$ (constant) and $Z_2 = \vec{u} \cdot \vec{X}$.

For positive constant a , $|\vec{w} + a\vec{u}| \leq |\vec{w}| + a|\vec{u}|$ which means that

$$\hat{\rho}[(\vec{w} + a\vec{u}) \cdot \vec{X}] + \eta|\vec{w} + a\vec{u}| \leq \hat{\rho}[\vec{w} \cdot X] + \eta|\vec{w}| - a(\mu - \eta)|\vec{u}| \quad (17)$$

For $\mu > \eta$ we can achieve portfolios with negative risk of arbitrarily large absolute value by letting $a \rightarrow \infty$: the solution of the optimization problem runs off to infinity. \square

Therefore, ℓ_1 regularization is not a good choice if the probability that a μ -dominant portfolio exists is non-zero. The probability that a μ -dominant portfolio exists depends on the distribution of returns, and it is therefore hard to make any general statements without making assumptions about this distribution. Yet, if the support of the distribution of returns extends to the whole real axis, then the probability that a μ -dominant portfolio exists is non-zero. In the next section we assume Gaussian (random) portfolios, providing some intuition about this probability.

Notice that the proof breaks down when the ℓ_1 norm is replaced by the ℓ_p norm with $p > 1$. Then the objective function

$$\hat{\rho}[\vec{w} \cdot \vec{X}] + \eta \sum_{i=1}^N |w_i|^p \sim \eta \sum_{i=1}^N |w_i|^p (1 + O(w^{-(p-1)})) \quad (18)$$

is dominated by the regularization term when the weights w_i diverge.

Conversely, for $p < 1$ it is the regularization term that becomes negligible compared to the risk measure. Hence regularization with $p < 1$ is not able to cure the instability, as soon as a μ -dominant portfolio arises, for any $\mu \geq 0$.

4 Regularized Expected Shortfall for large random portfolios.

In this section we consider the problem of RPO for the risk measure Expected Shortfall in the particular case where both the number of assets N and the number of observations T are large ($N, T \rightarrow \infty$), but their ratio is a finite constant. In order to derive quantitative results, we consider, for the sake of simplicity, the case of N i.i.d standard normal returns. We can extend the approach of [9] to include regularization (see also [44]). This provides us with typical results, i.e. results that occur with probability arbitrarily close to one in the limit $N, T \rightarrow \infty$.

For i.i.d. random returns, the true optimal solution is the one where all the weights are the same, $w_i = 1/N$, $\forall i$. Individual solutions based on specific samples will likely deviate from the true optimum, especially when the sample size, T , is not large enough. The optima for individual samples may deviate from symmetry, but averaging over the samples will necessarily restore it. Likewise, the ℓ_1 regularizer that tends to produce sparse solutions will not be able to make a distinction between the equivalent assets, therefore it will set some weights to zero at random for each sample. Our goal here is merely to demonstrate the effect of ℓ_1 on the instability of a portfolio optimized under ES. The real effect of ℓ_1 should be studied using a heterogeneous (non-i.i.d.) portfolio, a problem we leave for future research.

The optimization of historical Expected Shortfall can be formulated as a linear programming problem [12]. Expected Shortfall is a piecewise linear loss function, just like the Maximal Loss. The optimization problem, including the ℓ_p regularizer,

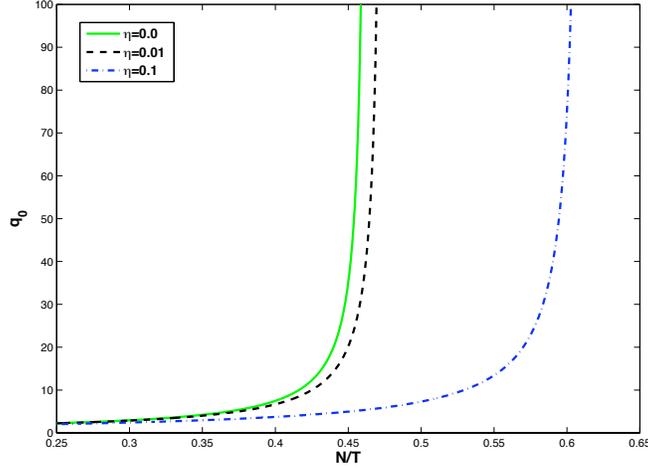


Figure 3: q_0 as a function of N/T for different values of η for the ℓ_1 regularized ES. The ℓ_1 regularizer shifts the feasible- infeasible transition towards larger values of N/T .

reads

$$\min_{\vec{w}, \vec{u}, \epsilon} \left[(1 - \beta)T\epsilon + \sum_{t=1}^T u_t + \eta |\vec{w}|^p \right], \quad (19)$$

$$\text{s.t.} \quad \vec{w} \cdot \vec{x}_t + \epsilon + u_t \geq 0; \quad u_t \geq 0; \quad \forall t, \quad (20)$$

$$\sum_i w_i = WN, \quad (21)$$

where W is a normalization parameter that is usually set to one. However, since market-impact depends on the size of the position that is liquidated, we explicitly keep track of W in what follows. In addition, the sum of the weights was chosen to be of the order of N , in order for the optimal weights to be of order one.

Assuming Gaussian returns, we can solve the above optimization problem analytically. Generalizing [9, 44] by including the regularizer, we follow lines familiar from the statistical physics of disordered systems: the loss function is regarded as the Hamiltonian (energy functional) of a fictitious statistical physics system, a fictitious temperature is introduced and the free energy of this system is calculated in the limit $N, T \rightarrow \infty$ with $T/N = \tau$ fixed. Averaging over the different realizations of the returns corresponds to what is called quenched averaging in statistical physics. We report here only the final result, and refer the interested reader to the Appendix, where details of the calculation are given. After some effort, one ends up with a free

energy functional, depending on six variables, the so-called order parameters:

$$\begin{aligned}
F(\lambda, \epsilon, q_0, \Delta, \hat{q}_0, \hat{\Delta}) &= \lambda W + \tau(1 - \beta)\epsilon - \Delta\hat{q}_0 - \hat{\Delta}q_0 \\
&+ \langle \min_w [V(w, z)] \rangle_z + \frac{\tau\Delta}{2\sqrt{\pi}} \int_{-\infty}^{\infty} ds e^{-s^2} g\left(\frac{\epsilon}{\Delta} + s\sqrt{\frac{2q_0}{\Delta^2}}\right),
\end{aligned} \tag{22}$$

where

$$V(w, z) = \hat{\Delta}w^2 + \eta|w|^p - \lambda w - zw\sqrt{-2\hat{q}_0}. \tag{23}$$

The function $g(x)$ is given in the Appendix. In Eq. (22) $\langle \cdot \rangle_z$ represents an average over the normal variable z . The “true” free energy, i.e. the minimal risk per asset, is the stationary point of the above expression with respect to variations of the order parameters $\lambda, \epsilon, q_0, \Delta, \hat{q}_0$ and $\hat{\Delta}$. The original optimization problem involving $N+T+1$ variables has been reduced, in the limit of large N and T , to the simpler problem of finding the minimum of the functional $F(\lambda, \epsilon, q_0, \Delta, \hat{q}_0, \hat{\Delta})$ that only depends on six variables.

4.1 ℓ_1 regularization

The origin of the instability for the Expected Shortfall under ℓ_1 is the same as with Maximal Loss: Expected Shortfall diverges (becomes unbounded from below) when there is a dominating portfolio [17]. The divergence of ES is linear, therefore a piecewise linear regularizer like ℓ_1 will not be able to eliminate the instability completely, it will do so only with a certain probability, depending on the sample. Thus we expect that after averaging over the ensemble of random returns the instability will persist. Yet, we know [9] that the probability to observe the instability is essentially zero for $N, T \rightarrow \infty$ when $\tau = T/N$ is larger than a critical threshold τ_c whereas it is arbitrarily close to one if $\tau < \tau_c$. We now illustrate how the instability arises by analyzing the extrema of Eq. (22), and show that ℓ_1 regularization merely shifts the instability to a smaller threshold.

The fact that the solution for $N, T \rightarrow \infty$ contains a nested optimization problem over w in the first term on the second line of Eq. (22) is not accidental. Indeed, this arises from the optimization over the weights w_i of the original problem. Therefore, the minimization of $V(w, z)$ can be thought of as a “representative weight” problem, where the random variable z over which the expectation is taken encodes the effect of the randomness in the sample.

The solution $w^*(z)$ of the representative weight problem

$$w^*(z) = \begin{cases} \frac{z\sqrt{-2\hat{q}_0} + \lambda - \eta}{2\hat{\Delta}}, & \text{if } z > \frac{\eta - \lambda}{\sqrt{-2\hat{q}_0}} \\ 0, & \text{if } \frac{-\eta - \lambda}{\sqrt{-2\hat{q}_0}} < z < \frac{\eta - \lambda}{\sqrt{-2\hat{q}_0}} \\ \frac{z\sqrt{-2\hat{q}_0} + \lambda + \eta}{2\hat{\Delta}}, & \text{if } z < \frac{-\eta - \lambda}{\sqrt{-2\hat{q}_0}}, \end{cases} \quad (24)$$

has the typical shape induced by ℓ_1 regularization. The Gaussian distribution of the z then induces the weights to be distributed according to $p(w) = \langle \delta(w - w^*(z)) \rangle_z$.

The first-order conditions for the free energy read

$$W = \langle w^* \rangle_z \quad (25)$$

$$(1 - \beta) + \frac{1}{2\sqrt{\pi}} \int_{-\infty}^{\infty} ds e^{-s^2} g' \left(\frac{\epsilon}{\Delta} + s\sqrt{\frac{2q_0}{\Delta^2}} \right) = 0 \quad (26)$$

$$\hat{\Delta} - \frac{\tau}{2\sqrt{2\pi q_0}} \int_{-\infty}^{\infty} ds e^{-s^2} s g' \left(\frac{\epsilon}{\Delta} + s\sqrt{\frac{2q_0}{\Delta^2}} \right) = 0 \quad (27)$$

$$-\hat{q}_0 - 2\frac{\hat{\Delta}q_0}{\Delta} + \frac{\tau}{2\sqrt{\pi}} \int_{-\infty}^{\infty} ds e^{-s^2} g \left(\frac{\epsilon}{\Delta} + s\sqrt{\frac{2q_0}{\Delta^2}} \right) + \tau(1 - \beta)\frac{\epsilon}{\Delta} = 0 \quad (28)$$

$$\Delta = \frac{1}{\sqrt{-2\hat{q}_0}} \langle w^* z \rangle_z \quad (29)$$

$$q_0 = \langle w^{*2} \rangle_z. \quad (30)$$

Notice in particular, that the first of these equations translates the budget constraint of the original problem into the properties of the representative weight problem above. The last equation, instead, relates the parameter $q_0 \propto \frac{1}{N} \sum_i w_i^2$ to the fluctuation of the weights. This corresponds to the distance of the solution from the optimal one (given by the constant portfolio). In the present setting, this is also related to the estimation error, since different samples will deviate from the optimal solution in different ways.

The instability of portfolio optimization manifests itself in the divergence of the parameter q_0 . This, in turn, implies that the distribution of weights $w^*(z)$ becomes degenerate, i.e. infinitely broad. Hence, some intuition about the origin of the instability can be gained from looking at the optimization problem for the representative weight. From the form of $V(w, z)$ we see that the optimization problem for w always

has a finite solution as long as the quantity $\hat{\Delta} > 0$. However, when $\hat{\Delta} = 0$, $V(w, z)$ can be unbounded from below. In this case, the estimated expected shortfall runs to minus infinity, because the term $\langle \min_w [V(w, z)] \rangle_z$ in the free energy diverges. Note that, because of the budget constraint implemented in equation (25), the optimal portfolio will in this case dictate to go infinitely long in some of the assets and infinitely short in some of the others. In summary, the condition $\hat{\Delta} = 0$ is what determines the critical point.

We numerically solved the system of equations (25)–(30), and we plot in Figure 3 the quantity q_0 as a function of the ratio $\tau^{-1} = N/T$ for different values of the regularizer’s amplitude η . The study of the stationarity conditions shows that, similarly to what we saw in the toy example, the ℓ_1 regularization produces a shift towards higher values of N/T of the feasible-infeasible transition, but does not eliminate it completely. The shift in the critical point where the estimation error blows up is clearly seen in Figure 3.

In the toy example discussed in section 3.1 we furthermore saw that the hard implementation of the ℓ_1 constraint, obtained in the limit $\eta \rightarrow \infty$, was equivalent to a ban on short selling. This is true also here. Indeed, if η goes to infinity, the regularizer imposes infinite penalty on any solution not on the ℓ_1 ball (defined by the relation $\sum_i |w_i| \leq NW$). On the other hand, the solution also has to be on the plane corresponding to the budget constraint. This means the solution must be on the $(+, +, + \dots +)$ face of the ℓ_1 ball, where all the weights are positive. Conversely, if all the weights are positive and they also satisfy the budget constraint, then the solution is necessarily on the $(+, +, + \dots +)$ face of the ℓ_1 ball.

That imposing a ban on short positions helps taming the large sample fluctuations was noticed by [24], though the authors did not make the connection to regularization. It is clear that any constraint that reduces the domain over which the optimum is sought to a finite volume has a similar effect: there cannot be infinite fluctuations in a finite volume.

4.2 ℓ_p -norm

There is no qualitative difference between the effects of the soft (finite η) or hard ($\eta \rightarrow \infty$) implementation of a non-linear regularizer that grows faster than linear, such as the ℓ_p -norm with $p > 1$; the regularizer will eliminate the instability either way.

We can see this by noting that if $p > 1$, then the function $V(w, z)$ always has a finite minimum for any value of z , thus preventing the divergence of the estimation error. The solution of the representative weight problem, w^* can be obtained by

minimization of the function $V(w, z)$ (Eq. 23). Let $p = 1 + \frac{1}{n}$. To compute w^* we must solve

$$2\hat{\Delta}w + \eta\left(1 + \frac{1}{n}\right)|w^*|^{1/n} - \lambda - z\sqrt{-2\hat{q}_0} = 0. \quad (31)$$

Taking the limit $\hat{\Delta} \rightarrow 0$, if $\eta > 0$ the equation reduces to

$$\eta\left(1 + \frac{1}{n}\right)|w^*|^{1/n} - \lambda - z\sqrt{-2\hat{q}_0} = 0, \quad (32)$$

which always has a finite solution for finite n . More generally, we can say that any regularizer that grows faster than linear will prevent the solutions from running away to infinity.

This is true in particular for the for the case $p = 2$ [44], and for the case $p = 3/2$, which corresponds to the square root market impact that characterizes the liquidation of large positions under normal market conditions [3–6].

4.3 Linear combination of the ℓ_1 and ℓ_2 regularizers.

In high-dimensional problems it is often expedient to use a combination of the ℓ_1 and ℓ_2 regularizers [7], sometimes called the elastic net [48]. This setting leads to the following optimization problem

$$\min_{\vec{w}, \vec{u}, \epsilon} \left[(1 - \beta)T\epsilon + \sum_{t=1}^T u_t + \eta_1|\vec{w}| + \eta_2|\vec{w}|^2 \right], \quad (33)$$

$$\text{s.t.} \quad \vec{w} \cdot \vec{x}_\tau + \epsilon + u_t \geq 0; \quad u_t \geq 0; \quad \forall t, \quad (34)$$

$$\sum_i w_i = WN. \quad (35)$$

The replica method readily extends also to this case. In this situation the free energy is again of the form (22), but with

$$V(w, z) = \hat{\Delta}w^2 + \eta_1|w| + \eta_2|w|^2 - \lambda w - zw\sqrt{-2\hat{q}_0}. \quad (36)$$

The solution of the optimization problem for the representative weight reads:

$$w^*(z) = \begin{cases} \frac{z\sqrt{-2\hat{q}_0} + \lambda - \eta_1}{2(\hat{\Delta} + \eta_2)}, & \text{if } z > \frac{\eta_1 - \lambda}{\sqrt{-2\hat{q}_0}} \\ 0, & \text{if } \frac{-\eta_1 - \lambda}{\sqrt{-2\hat{q}_0}} < z < \frac{\eta_1 - \lambda}{\sqrt{-2\hat{q}_0}} \\ \frac{z\sqrt{-2\hat{q}_0} + \lambda + \eta_1}{2(\hat{\Delta} + \eta_2)}, & \text{if } z < -\frac{\eta_1 + \lambda}{\sqrt{-2\hat{q}_0}}. \end{cases} \quad (37)$$

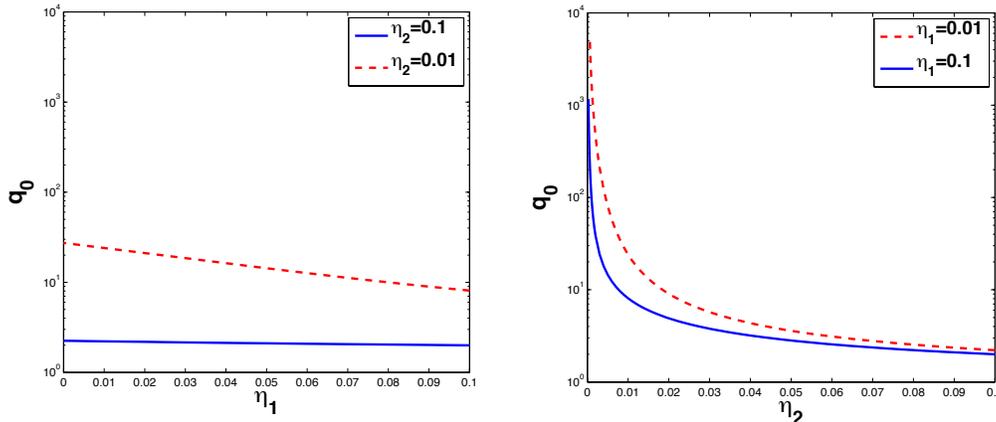


Figure 4: **Left panel:** q_0 as a function of η_1 for $\eta_2 = 0.1$ (blue solid line) and $\eta_2 = 0.01$ (red dashed line). **Right panel:** q_0 as a function of η_2 for $\eta_1 = 0.1$ (blue solid line) and $\eta_1 = 0.01$ (red dashed line). In both panels $\tau = 1.5$. The quantity q_0 strongly depends on η_2 , in particular close to $\eta_2 = 0$.

From these expressions we see that the divergence of w^* is prevented as long as $\eta_2 > 0$, while a divergence may occur if $\eta_2 = 0$ when $\hat{\Delta} = 0$.

Using this form of $w^*(z)$ to evaluate the averages over the variable z that appear in equations (25), (29) and (30), we numerically solve the set of first-order conditions (25)–(30).

We show the behavior of q_0 as a function of the parameters η_1 and η_2 , for $\tau = 1.5$, in Figure 4. The dependence on η_2 is naturally stronger than the dependence on η_1 .

A useful scheme may be to set η_2 to a value that ensures small estimation error given the value of τ in a particular application, and then to vary η_1 in order to achieve the desired level of sparseness. Although sparseness reduces diversification, there might be situations where limiting the dimension of a portfolio may be advantageous, for example by limiting transaction costs.

5 Conclusion

We have shown that market impact considerations will drive an investor to Regularized Portfolio Optimization (RPO) under any coherent risk measure. The impact function determines which regularizer should be used. Generalizing our previous result, that linear impact leads to the use of the ℓ_2 norm regularizer [44], we have

shown here how the impact function selects the regularizer in general.

We have also demonstrated that the instability of Expected Shortfall found in [9] will be removed by any nonlinear regularizer that grows faster than linear, but will only be shifted by the piecewise linear ℓ_1 norm based regularizer that corresponds to a bid-ask spread impact. Although we wanted to mostly concentrate on Expected Shortfall in this paper, we would like to point out that due to the positive homogeneity and translational invariance axioms, ℓ_p regularizers will have the same effect on any coherent measure as they do on ES. (Moreover, as these properties are shared also by Value at Risk, this extends also to VaR.)

In order to illustrate the above points, we also provided explicit analytical calculations for the optimization of large random portfolios. These calculations required the use of techniques borrowed from the statistical physics of random systems where they had been widely used and rigorously established in the cases of various objective functions [49–52]. As these techniques may not be familiar for most of our readers, we tried to simplify the task as much as possible, and constrained ourselves to the case of i.i.d. normal returns. This is clearly not realistic, as returns in real markets exhibit heavy tails and display complex patterns of interdependencies (see for instance [53]). The assumption of independence is not crucial, correlated normal variables can be easily accommodated by the formalism and do not change the conclusions qualitatively, as has been shown in a different example by [16]. So our analysis could be extended analytically to the case of multi-factor models, to account for correlations between variables due, for example, to industrial sectors or geographic regions.

As for the Gaussian character, it is a critical element of the formalism. In order to tackle fat tailed distributions, one has to resort to numerical simulations, but it is a foregone conclusion again that they cannot qualitatively modify the main message, because the correspondence between the character of the market impact and the corresponding regularizer follows directly from the coherence axioms and is obviously independent of the statistical properties of the returns, while the claim about which regularizer can take care of the instability of which risk measure depends only on the asymptotic behavior of these quantities as functions of the weights, not on the returns.

Another simplifying feature of the illustrative calculation is the use of the large N limit. This is a technicality: at a certain point we had to use the saddle-point method to evaluate some integrals. The quality of the saddle-point approximation depends on the curvature of the cost function around the saddle-point, and varies from point to point in the plane of the control parameters N/T and the confidence limit β . General experience with saddle point integrals (think of the Stirling-formula) supported by numerical simulations in our present case tells us that the N values

required for a good approximation can be quite moderate (such as 50 to 100) if we are safely inside the stable region, while in the immediate vicinity of the instability line much larger values (500 or above) are needed.

In summary, we believe that the qualitative message of our paper is robust: The instability of Expected Shortfall is removed by any ℓ_p regularizer with $p > 1$ and shifted by an ℓ_1 regularizer.

We find it appealing that the statistically motivated use of regularizers is intimately linked to market impact: consideration of finite liquidity should limit excessively large positions, which is exactly what regularizers do.

We also find it remarkable how powerful the method of replicas proves to be in producing analytical results for the estimation error in situations which would probably be very hard to approach by more conventional methods of probability theory.

Acknowledgement

I.K. has been supported by the European Union under grant agreement No. FP7-ICT-255987-FOC-II Project and by the Institute for New Economic Thinking under grant agreement ID: INO1200019. S.S. thanks the Abdus Salam International Center for Theoretical Physics for their hospitality and support of this collaboration. M.M. acknowledges support from the Marie Curie Training Network NETADIS (FP7 - Grant 290038). F.C. acknowledges support of the Economic and Social Research Council (ESRC) in funding the Systemic Risk Centre (ES/K002309/1).

Appendix

In this appendix we show how, in the limit where N and T go to infinity, the optimization problem defined by equations (19), (20) and (21) can be reduced to that of finding the minimum of the “free energy” functional given by (22).

We need to find the minimum of

$$E[\epsilon, \{u_t\}] = (1 - \beta)T\epsilon + \sum_{t=1}^T u_t + \eta \|\vec{w}\|_p$$

under the constraints

$$u_t \geq 0, \quad u_t + \epsilon + \sum_{i=1}^N x_{i,t} w_i \geq 0 \quad \forall t$$

and

$$\sum_{i=1}^N w_i = WN.$$

The calculation proceeds as follows: Given a history of returns $\{x_{i,t}\}$, we introduce the inverse temperature parameter γ and define the canonical partition function (or generating functional) as

$$Z_\gamma[\{x_{i,t}\}] = \int_0^\infty \prod_{i=1}^T du_t \int_{-\infty}^\infty d\epsilon \theta\left(u_t + \epsilon + \sum_{i=1}^N x_{i,t} w_i\right) \delta\left(\sum_{i=1}^N w_i - WN\right) e^{-\gamma E[\epsilon, \{u_t\}]}, \quad (38)$$

where $\theta(x) = 1$ if $x > 0$ and zero otherwise. The partition function is therefore an integral over all possible configurations of variables that are compatible with the constraints of the problem, where each configuration $\epsilon, \{u_t\}$ is weighted by the quantity $e^{-\gamma E[\epsilon, \{u_t\}]}$ (the Boltzmann weight). For large values of γ the Boltzmann weight is concentrated on the configurations corresponding to the smallest values of $E[\epsilon, \{u_t\}]$, so the minimum of the free energy (or cost function) can be obtained by letting the inverse temperature go to infinity: The minimum cost (per asset) can then be computed in the limit of large N as

$$\lim_{N \rightarrow \infty} \lim_{\gamma \rightarrow \infty} -\frac{\log Z_\gamma[\{x_{i,t}\}]}{\gamma N}. \quad (39)$$

To derive the typical properties of the ensemble we have to average over all possible realizations of returns and compute

$$\langle \log Z_\gamma[\{x_{i,t}\}] \rangle = \int_{-\infty}^\infty \prod_{i=1}^N \prod_{t=1}^T dx_{i,t} P[\{x_{i,t}\}] \log Z_\gamma[\{x_{i,t}\}], \quad (40)$$

where $P[\{x_{i,t}\}]$ is the probability density function of returns. The calculation of the average of a logarithm is hard. To circumvent this difficulty, it is expedient to apply the so called “replica trick”, a device that dates back to [54] and has since been used in various problems in the statistical physics of random systems, most notably in spin glass theory [49]. The “trick” consist in using the identity

$$\langle \log Z \rangle = \lim_{n \rightarrow 0} \frac{\partial \langle Z^n \rangle}{\partial n}. \quad (41)$$

For integer n , we can compute the moment $\langle Z^n \rangle$ as the average partition function of a system composed of n independent replicas of the original one - hence the name

of the trick. An analytical continuation to real values of n will then allow us to perform the limit $n \rightarrow 0$ and obtain the sought quantity $\langle \log Z_\gamma [\{x_{i,t}\}] \rangle$. The astute reader notices that the procedure involves some suspicious steps (order of limits, non-uniqueness of the analytic continuation from the natural numbers to reals). These have been carefully analyzed and the results obtained via the use of the trick reproduced by rigorous methods [50, 51], [52] in problems of a similar nature (with convex cost functions), but also in the much more complicated case where the cost function had a large number of competing minima. For the special problem of ES treated here no such rigorous proof is known, so the replica trick can be regarded as a heuristic method that produces highly nontrivial results. These can a posteriori be verified by simulations. We will provide some simulation results as an illustration at the end of this appendix. Returning to the main line of argument, the replicated partition function can be computed as⁴

$$\begin{aligned}
Z_\gamma^n &= \int_{-\infty}^{\infty} \left(\prod_{i=1}^N \prod_{t=1}^T dx_{i,t} \right) \int_{-\infty}^{\infty} \left(\prod_{a=1}^n d\epsilon^a \right) \int_0^{\infty} \left(\prod_{t=1}^T \prod_{a=1}^n du_t^a \right) \int_{-\infty}^{\infty} \left(\prod_{i=1}^N \prod_{a=1}^n dw_i^a \right) \\
&\times \int_{-i\infty}^{i\infty} \left(\prod_{a=1}^n d\hat{\lambda}^a \right) \int_0^{\infty} \left(\prod_{t=1}^T \prod_{a=1}^n d\mu_t^a \right) \int_{-\infty}^{\infty} \left(\prod_{t=1}^T \prod_{a=1}^n d\hat{\mu}_t^a \right) \prod_{t=1}^T \prod_{i=1}^N \exp \left\{ -\frac{Nx_{i,t}^2}{2} \right\} \\
&\times \exp \left\{ \sum_a \hat{\lambda}^a \left(\sum_i w_i^a - WN \right) \right\} \prod_t \exp \left\{ \sum_a i\hat{\mu}_t^a \left(u_t^a + \epsilon^a + \sum_i x_{i,t} w_i^a - \mu_t^a \right) \right\} \\
&\times \exp \left\{ -\gamma \sum_a (1-\beta) T \epsilon^a - \gamma \sum_{a,t} u_t^a - \gamma \eta \sum_i |w_i^a|^p \right\},
\end{aligned}$$

where we have assumed that

$$P[\{x_{i,t}\}] = \prod_{t=1}^T \prod_{i=1}^N \exp \left\{ -\frac{Nx_{i,t}^2}{2} \right\}, \quad (42)$$

and we have enforced the constraints through the Lagrange multipliers $\hat{\lambda}^a$, μ_t^a and $\hat{\mu}_t^a$. Averaging over the quenched variables $\{x_{i,t}\}$ and introducing the overlap matrix $Q_{a,b} = \frac{1}{N} \sum_i w_i^a w_i^b$ and its conjugate $\hat{Q}_{a,b}$ one obtains

⁴In the calculation we will not keep track of constant multiplicative factors that do not affect the final result.

$$\begin{aligned}
Z_\gamma^n &= \int_{-i\infty}^{i\infty} \left(\prod_{a=1}^n \prod_{b=1}^n dQ_{a,b} d\hat{Q}_{a,b} \right) \int_{-\infty}^{\infty} \left(\prod_{a=1}^n d\epsilon^a \right) \int_0^\infty \left(\prod_{t=1}^T \prod_{a=1}^n du_t^a \right) \int_{-\infty}^{\infty} \left(\prod_{i=1}^N \prod_{a=1}^n dw_i^a \right) \\
&\times \int_{-i\infty}^{i\infty} \left(\prod_{a=1}^n d\hat{\lambda}^a \right) \int_0^\infty \left(\prod_{t=1}^T \prod_{a=1}^n d\mu_t^a \right) \int_{-\infty}^{\infty} \left(\prod_{t=1}^T \prod_{a=1}^n d\hat{\mu}_t^a \right) \exp \left\{ \sum_a \hat{\lambda}^a \left(\sum_i w_i^a - WN \right) \right\} \\
&\times \prod_t \exp \left\{ -\frac{1}{2} \sum_{a,b} \hat{\mu}_t^a Q_{a,b} \hat{\mu}_t^b \right\} \exp \left\{ \sum_{a,b} \hat{Q}_{a,b} \left(NQ_{a,b} - \sum_i w_i^a w_i^b \right) \right\} \\
&\times \exp \left\{ -\gamma \sum_a (1-\beta) T \epsilon^a - \gamma \sum_{a,t} u_t^a - \gamma \eta \sum_i |w_i^a|^p \right\} \\
&\times \prod_t \exp \left\{ i \sum_a \hat{\mu}_t^a (u_t^a + \epsilon^a - \mu_t^a) \right\}.
\end{aligned}$$

We can now perform the Gaussian integral over the variables $\{\hat{\mu}_t^a\}$:

$$\begin{aligned}
Z_\gamma^n &= \int_{-i\infty}^{i\infty} \left(\prod_{a=1}^n \prod_{b=1}^n dQ_{a,b} d\hat{Q}_{a,b} \right) \int_{-\infty}^{\infty} \left(\prod_{a=1}^n d\epsilon^a \right) \int_0^\infty \left(\prod_{t=1}^T \prod_{a=1}^n du_t^a \right) \int_{-\infty}^{\infty} \left(\prod_{i=1}^N \prod_{a=1}^n dw_i^a \right) \\
&\times \int_{-i\infty}^{i\infty} \left(\prod_{a=1}^n d\hat{\lambda}^a \right) \int_0^\infty \left(\prod_{t=1}^T \prod_{a=1}^n d\mu_t^a \right) \exp \left\{ \sum_a \hat{\lambda}^a \left(\sum_i w_i^a - WN \right) \right\} \\
&\times \exp \left\{ -\gamma \sum_a (1-\beta) T \epsilon^a - \gamma \sum_{a,t} u_t^a - \gamma \eta \sum_i |w_i^a|^p \right\} \exp \left\{ \sum_{a,b} \hat{Q}_{a,b} \left(NQ_{a,b} - \sum_i w_i^a w_i^b \right) \right\} \\
&\times \prod_t \exp \left\{ -\frac{1}{2} \sum_{a,b} (u_t^a + \epsilon^a - \mu_t^a) Q_{a,b}^{-1} (u_t^b + \epsilon^b - \mu_t^b) \right\} \exp \left\{ -\frac{T}{2} \text{tr} \log Q \right\}.
\end{aligned}$$

Introducing the variables $y_t^a = \mu_t^a - u_t^b$ and $z_t^a = \mu_t^a + u_t^b$ and integrating over the $\{z_t^a\}$ we obtain

$$\begin{aligned}
Z_\gamma^n &= \int_{-i\infty}^{i\infty} \left(\prod_{a=1}^n \prod_{b=1}^n dQ_{a,b} d\hat{Q}_{a,b} \right) \int_{-\infty}^{\infty} \left(\prod_{a=1}^n d\epsilon^a \right) \int_{-\infty}^{\infty} \left(\prod_{i=1}^N \prod_{a=1}^n dw_i^a \right) \int_{-i\infty}^{i\infty} \left(\prod_{a=1}^n d\hat{\lambda}^a \right) \\
&\times \exp \left\{ \sum_a \hat{\lambda}^a \left(\sum_i w_i^a - WN \right) \right\} \exp \left\{ -\gamma \sum_a (1-\beta) T \epsilon^a - \gamma \sum_{a,t} u_t^a - \gamma \eta \sum_i |w_i^a|^p \right\} \\
&\times \exp \left\{ \sum_{a,b} \hat{Q}_{a,b} \left(N Q_{a,b} - \sum_i w_i^a w_i^b \right) \right\} \prod_t \exp \left\{ -\frac{1}{2} \sum_{a,b} (u_t^a + \epsilon^a - \mu_t^a) Q_{a,b}^{-1} (u_t^b + \epsilon^b - \mu_t^b) \right\} \\
&\times \exp \left\{ -\frac{T}{2} \text{tr} \log Q - TN \log \gamma + T \log Z_\gamma(\{\epsilon^a, Q\}) \right\}
\end{aligned}$$

where

$$\begin{aligned}
Z_\gamma(\{\epsilon^a, Q\}) &= \int_{-\infty}^{+\infty} \prod_a dy^a \exp \left\{ -\frac{1}{2} \sum_{a,b} (y^a - \epsilon^a) Q_{a,b}^{-1} (y^b - \epsilon^b) \right\} \\
&\times \exp \left\{ \gamma \sum_a y^a \theta(-y^a) \right\}.
\end{aligned}$$

At this point we have to make a guess concerning the structure of the overlap matrices $Q_{a,b}$ and $\hat{Q}_{a,b}$. The symmetry of the problem with respect to the permutation of replicas leads one to the “replica symmetric” ansatz

$$Q_{a,b} = \begin{cases} q_1, & a = b \\ q_0, & a \neq b \end{cases} \quad (43)$$

$$\hat{Q}_{a,b} = \begin{cases} r_1, & a = b \\ r_0, & a \neq b. \end{cases} \quad (44)$$

The assumption of replica symmetry has been justified in every earlier application of the trick where the cost function was convex.

The next step is to perform the “zero temperature” limit $\gamma \rightarrow \text{infy}$. This can be effected by introducing the following rescaling relations

$$\Delta = \gamma(q_1 - q_0), \quad (45)$$

$$\hat{\Delta} = (r_1 - r_0)/\gamma, \quad (46)$$

$$\lambda^a = \hat{\lambda}^a \gamma, \quad (47)$$

$$\hat{q}_0 = r_0 \gamma^2. \quad (48)$$

The \vec{w} -dependent part of the partition function is

$$\int [Dw] e^{-\gamma F_w} = \int [Dw] e^{\sum_{ia} \lambda^a w_i^a - \gamma \eta \sum_i |w_i^a|^p - \sum_{a,b} \hat{Q}_{a,b} \sum_i w_i^a w_i^b} \quad . \quad (49)$$

Exploiting the identity $\log \langle X^n \rangle \simeq n \langle \log X \rangle$ valid for $n \rightarrow 0$, and after some manipulations, we arrive at the following contribution to the free energy

$$F_w = \gamma \left\langle \log \int dw e^{-\gamma [\hat{\Delta} w^2 + \eta |w|^p - \lambda w - zw \sqrt{-2\hat{q}_0}]} \right\rangle_z, \quad (50)$$

where the notation $\langle \dots \rangle_z$ means averaging over the normal variable z . After some further manipulations, we can write the partition function as

$$Z_\gamma^n = \int d\lambda d\epsilon dq_0 d\Delta d\hat{q}_0 d\hat{\Delta} e^{-\gamma n N F[\lambda, \epsilon, q_0, \Delta, \hat{q}_0, \hat{\Delta}]} \quad (51)$$

where

$$\begin{aligned} F(\lambda, \epsilon, q_0, \Delta, \hat{q}_0, \hat{\Delta}) &= \lambda W + \tau(1 - \beta)\epsilon - \Delta \hat{q}_0 - \hat{\Delta} q_0 \quad (52) \\ &- \frac{1}{\gamma} \left\langle \log \int_{-\infty}^{\infty} dw e^{-\gamma V(w, z)} \right\rangle_z + \frac{\tau \Delta}{2\sqrt{\pi}} \int_{-\infty}^{\infty} ds e^{-s^2} g \left(\frac{\epsilon}{\Delta} + s \sqrt{2 \frac{q_0}{\Delta^2}} \right), \end{aligned}$$

with $\tau = T/N$,

$$V(w, z) = \hat{\Delta} w^2 + \eta |w|^p - \lambda w - zw \sqrt{-2\hat{q}_0} \quad (53)$$

and

$$g(x) = \begin{cases} 0, & x \geq 0 \\ x^2, & -1 \leq x \leq 0 \\ -2x - 1, & x < -1 \end{cases} \quad . \quad (54)$$

In the limit of large N the integral in equation (51) is concentrated around the minimum of $F(\lambda, \epsilon, q_0, \Delta, \hat{q}_0, \hat{\Delta})$ and can be computed through the saddle point method.

The minimum cost

$$\lim_{\gamma \rightarrow \infty} \lim_{n \rightarrow 0} \left(-\frac{1}{\gamma} \frac{\partial Z_\gamma^n}{\partial n} \right) \quad (55)$$

corresponds then to the minimum of the free energy functional $F(\lambda, \epsilon, q_0, \Delta, \hat{q}_0, \hat{\Delta})$ in the limit $\gamma \rightarrow \infty$.

The saddle point method produces an asymptotic series. The corrections to the leading term given above depend on the curvature of the exponent in the integral.

These corrections are rather small away from the critical point, and the results obtained for asymptotically large N give a fairly good representation of the behaviour of the simulations for finite, moderately large N values, such as 50 to a few hundred. This is clearly shown in figure 5, where we present a comparison between the analytical solution obtained through the replica method (the lines in the plot) and results from a numerical solution of the optimization problem for $N = 128$ assets for the case of an ℓ_1 regularizer. We see from the plot that the agreement between the two is quite good already for this relatively small value of N .

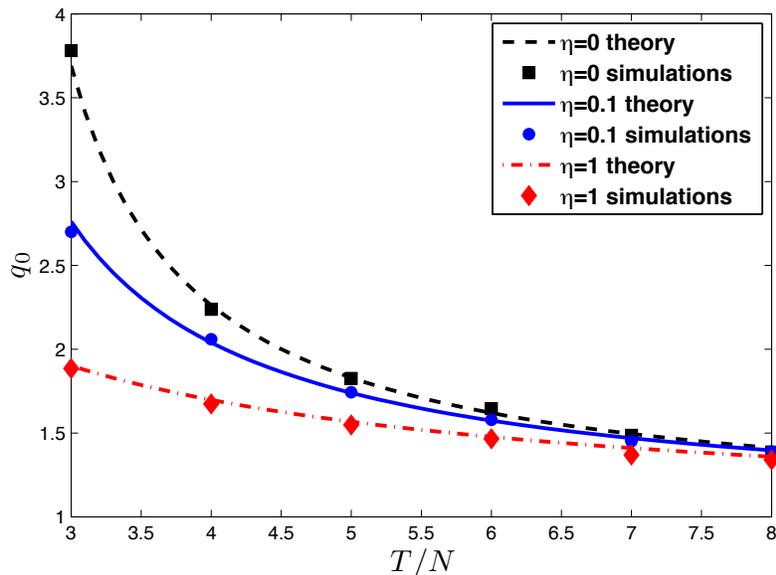


Figure 5: Comparison between analytical solution and numerical simulations. q_0 as a function of the ratio T/N for the case of an ℓ_1 regularizer and different values of η . Numerical simulations refer to systems with $N = 128$ assets. The match between the result of the replica calculation and numerical simulations is very good in this regime.

It has to be noted however that in the immediate vicinity of the critical point much larger N values should be used in the simulations in order to obtain agreement with the analytical calculation, as we show in figure 6.

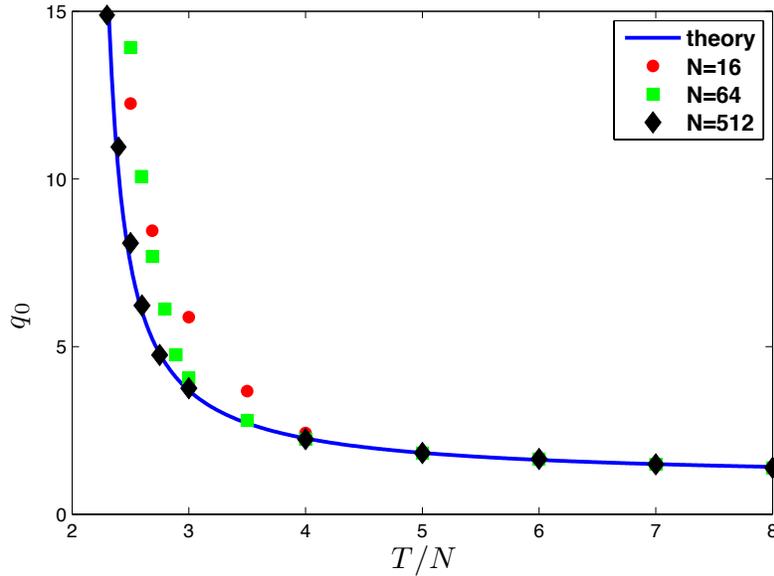


Figure 6: Comparison between analytical solution and numerical simulations. q_0 as a function of the ratio T/N for $\eta = 0$. Numerical simulations refer to systems with $N = 16, 64, 512$ assets. As we get closer to the critical point (which in this case is for T/N around 2.2), finite size effects become apparent and larger values of N are needed to obtain a good match with the analytical curve.

References

- [1] Carlo Acerbi and Giacomo Scandolo. Liquidity risk theory and coherent measures of risk. *Quantitative Finance*, 8(7):681–692, 2008.
- [2] J-P Bouchaud. The endogenous dynamics of markets: Price impact, feedback loops and instabilities. In A. Berd, editor, *Lessons from the 2008 crisis*, pages 57–156. Risk publications, 2011.
- [3] B. Toth, Y. Lempriere, C. Deremble, J. de Lataillade, J. Kockelkoren, and J-P. Bouchaud. Anomalous price impact and the critical nature of liquidity in financial markets. *Physical Review X*, 1(2):021006, 2011. <http://arxiv.org/abs/1105.1694>.

- [4] E. Moro, L. G. Moyano, J. Vicente, A. Gerig, J. D. Farmer, G. Vaglica, F. Lillo, and R.N. Mantegna. Market impact and trading protocols of hidden orders in stock markets. *Physical Review E*, 80(6):066102, 2009.
- [5] J. D. Farmer, A. Gerig, F. Lillo, and H. Waelbroeck. How efficiency shapes market impact. *Quantitative Finance*, 13(11):1743–1758, 2013.
- [6] F. Caccioli, J.-P. Bouchaud, and Farmer J. D. A proposal for impact-adjusted valuation: Critical leverage and execution risk. *arXiv:1204.0922*, 2012.
- [7] P. Bühlmann and S. van de Geer. *Statistics for High-Dimensional Data: Methods, Theory and Applications*. Springer Series in Statistics. Springer, Berlin, Heidelberg, 2011.
- [8] I. Kondor, S. Pafka, and G. Nagy. Noise sensitivity of portfolio selection under various risk measures. *Journal of Banking and Finance*, 31:1545–1573, 2007.
- [9] S. Ciliberti, I. Kondor, and M. Mezard. On the feasibility of portfolio optimization under expected shortfall. *Quantitative Finance*, 7:389–396, 2007.
- [10] M. Mézard and A. Montanari. *Information, Physics, and Computation*. Oxford Graduate Texts. OUP Oxford, 2009.
- [11] I. Varga-Haszonits and I. Kondor. Noise Sensitivity of Portfolio Selection in Constant Conditional Correlation GARCH models. *Physica*, A385:307–318, 2007.
- [12] R. T. Rockafellar and S. Uryasev. Optimization of conditional value-at-risk. *Journal of Risk*, 2(3):21–41, 2000.
- [13] C. Acerbi, C. Nordio, and C. Sirtori. Expected Shortfall as a Tool for Financial Risk Management, February 2001.
- [14] Basel Committee on Banking Supervision. Fundamental review of the trading book, 2012.
- [15] Basel Committee on Banking Supervision . Fundamental review of the trading book: A revised market risk framework., 2013.
- [16] I. Varga-Haszonits and I. Kondor. The instability of downside risk measures. *J. Stat. Mech.*, P12007, 2008.
- [17] I. Kondor and I. Varga-Haszonits. Instability of portfolio optimization under coherent risk measures. *Advances in Complex Systems*, 13, 2010.

- [18] P. Artzner, F. Delbaen, J.-M. Eber, and D. Heath. Coherent Measures of Risk. *Mathematical Finance*, 9:203–228, 1999.
- [19] E. J. Elton and M. J. Gruber. *Modern Portfolio Theory and Investment Analysis*. Wiley, New York, 1995.
- [20] J.D. Jobson and B. Korkie. Improved Estimation for Markowitz Portfolios Using James-Stein Type Estimators. *Proceedings of the American Statistical Association (Business and Economic Statistics)*, 1:279–284, 1979.
- [21] P. Jorion. Bayes-Stein Estimation for Portfolio Analysis. *Journal of Financial and Quantitative Analysis*, 21:279–292, 1986.
- [22] P.A. Frost and J.E. Savarino. An Empirical Bayes Approach to Efficient Portfolio Selection. *Journal of Financial and Quantitative Analysis*, 21:293–305, 1986.
- [23] R. Macrae and C. Watkins. Safe portfolio optimization. In H. Bacelar-Nicolau, F. C. Nicolau, and J. Janssen, editors, *Proceedings of the IX International Symposium of Applied Stochastic Models and Data Analysis: Quantitative Methods in Business and Industry Society, ASMDA-99, 14-17 June 1999, Lisbon, Portugal*, page 435. INE, Statistics National Institute, Portugal, 1999.
- [24] R. Jagannathan and T. Ma. Risk reduction in large portfolios: Why imposing the wrong constraints helps. *Journal of Finance*, 58:1651–1684, 2003.
- [25] V. DeMiguel, L. Garlappi, and R. Uppal. Optimal versus Naive Diversification: How Efficient is the 1/N Portfolio Strategy? *Review of Financial Studies*, 2007.
- [26] L. Garlappi, R. Uppal, and T. Wang. Portfolio selection with parameter and model uncertainty: a multi-prior approach. *Review of Financial Studies*, 20:41–81, 2007.
- [27] R. Kan and G. Zhou. Optimal portfolio choice with parameter uncertainty. *Journal of Financial and Quantitative Analysis*, 42:621–656, 2007.
- [28] G. Frahm and Ch. Memmel. Dominating estimators for the global minimum variance portfolio, 2009. Deutsche Bundesbank, Discussion Paper, Series 2: Banking and Financial Studies.
- [29] L. Laloux, P. Cizeau, J.-P. Bouchaud, and M. Potters. Noise Dressing of Financial Correlation Matrices. *Phys. Rev. Lett.*, 83:1467–1470, 1999.

- [30] 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, 1999.
- [31] L. Laloux, P. Cizeau, J.-P. Bouchaud, and M. Potters. Random Matrix Theory and Financial Correlations. *International Journal of Theoretical and Applied Finance*, 3:391, 2000.
- [32] V. Plerou, P. Gopikrishnan, B. Rosenow, L. A. N. Amaral, and H. E. Stanley. Universal and nonuniversal properties of cross correlations in financial time series. *Phys. Rev. Lett.*, 83:1471–1474, Aug 1999.
- [33] Z. Burda, A. Goerlich, and A. Jarosz. Signal and noise in correlation matrix. *Physica*, A343:295, 2004.
- [34] M. Potters and J.-Ph. Bouchaud. Financial applications of random matrix theory: Old laces and new pieces. *Acta Phys. Pol.*, B36:2767, 2005.
- [35] O. Ledoit and M. Wolf. Improved Estimation of the Covariance Matrix of Stock Returns with an Application to Portfolio Selection. *Journal of Empirical Finance*, 10(5):603–621, 2003.
- [36] O. Ledoit and M. Wolf. A well-conditioned estimator for large-dimensional covariance matrices. *J. Multivar. Anal.*, 88:365–411, 2004.
- [37] O. Ledoit and M. Wolf. Honey, I Shrunk the Sample Covariance Matrix. *J. Portfolio Management*, 31:110, 2004.
- [38] V. Golosnoy and Y. Okhrin. Multivariate shrinkage for optimal portfolio weights. *The European Journal of Finance*, 13:441–458, 2007.
- [39] V. DeMiguel, L. Garlappi, F. J. Nogales, and R. Uppal. A generalized approach to portfolio optimization: Improving performance by constraining portfolio norms. *Management Science*, 55:798–812, 2009.
- [40] J. Brodie, I. Daubechies, C. De Mol, D. Giannone, and I. Loris. Sparse and stable markowitz portfolios. *Proceedings of the National Academy of Sciences*, 106(30):12267–12272, 2009.
- [41] S. Still and I. Kondor. Regularizing portfolio optimization. *New Journal of Physics*, 12(7):075034, 2010.

- [42] A. Takeda and M. Sugiyama. ν -support vector machine as conditional value-at-risk minimization. In A. McCallum and S. Roweis, editors, *Proceedings of the 25th international conference on Machine learning (ICML)*, volume 307, pages 1056–1063. Omnipress, 2008.
- [43] B. Schölkopf, A. J. Smola, R. C. Williamson, and P. L. Bartlett. New support vector algorithms. *Neural Computation*, 12(5):1207–1245, 05 2000.
- [44] F. Caccioli, S. Still, M. Marsili, and I. Kondor. Optimal liquidation strategies regularize portfolio selection. *The European Journal of Finance*, 19(6):554–571, 2013.
- [45] Albert S Kyle and Anna A Obizhaeva. Market microstructure invariants: empirical evidence from portfolio transitions. *Available at SSRN 1978943*, 2011.
- [46] Albert S Kyle and Anna A Obizhaeva. Market microstructure invariance: Theory and empirical tests. Technical report, working paper, University of Maryland, 2013.
- [47] M. R. Young. A minimax portfolio selection rule with linear programming solution. *Management Science*, 44:673–683, 1998.
- [48] Hui Zou and Trevor Hastie. Regularization and variable selection via the elastic net. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 67(2):301–320, 2005.
- [49] M. Mézard, G. Parisi, and M. A. Virasoro. *Spin glass theory and beyond*. World Scientific Lecture Notes in Physics Vol. 9, World Scientific, Singapore, 1987.
- [50] F. Guerra and F. L. Toninelli. The infinite volume limit in generalized mean field disordered models. *Markov Proc. Rel. Fields*, 9:195–207, 2003.
- [51] F. Guerra and F. L. Toninelli. The thermodynamic limit in mean field spin glass model. *Commun. Math. Phys.*, 230:71–79, 2002.
- [52] M. Talagrand. *Spin glasses: a challenge for mathematicians: cavity and mean field models*, volume 46. Springer Science & Business Media, 2003.
- [53] Dror Y Kenett, Xuqing Huang, Irena Vodenska, Shlomo Havlin, and H Eugene Stanley. Partial correlation analysis: applications for financial markets. *Quantitative Finance*, 15(4):569–578, 2015.

- [54] Godfrey Harold Hardy, John Edensor Littlewood, and George Pólya. *Inequalities*. Cambridge university press, 1952.