Relative Robust Portfolio Optimisation with Benchmark Regret

Gonçalo Simões*, Mark McDonald†, Stacy Williams, Daniel Fenn‡ and Raphael Hauser§

March 20, 2018

Abstract

We extend Relative Robust Portfolio Optimisation models to allow portfolios to optimise their perfor- mance when considered relative to a set of benchmarks. We do this in a minimum volatility setting, where we model regret directly as the maximum different between our volatility and that of a given benchmark. Portfolio managers are also given the option of computing regret as a proportion of the benchmarks performance, which is more in line with market practice than other approaches suggested in the litera- ture. Furthermore we propose using regret as an extra constraint rather than as a brand new objective function, so practitioners can maintain their current framework. We also look into how such a triple optimisation problem can be solved or at least approximated for a general class of objective functions and uncertainty and benchmark sets. Finally we illustrate the benefits of this approach by examining its performance against other common methods in the literature in several equity markets.

1 Introduction

Modern Portfolio Theory has been an area of active research in mathematical finance since [Mar52], but it has not been fully adopted by practitioners yet.

*Corresponding author: Mathematical Institute, Radcliffe Observatory Quarter, Woodstock Road, Oxford OX2 6GG, U.K.; simoes@maths.ox.ac.uk; Oxford-Man Institute of Quantitative Finance, University of Oxford, Oxford, OX2 6ED, UK

†FX Quantitative Research, HSBC Bank, London E14 5HQ, UK

‡FX Quantitative Research, HSBC Bank, London E14 5HQ, UK

§Mathematical Institute, Radcliffe Observatory Quarter, Woodstock Road, Oxford OX2 6GG, U.K.; Pembroke College, St Aldates, Oxford OX1 1DW, UK

One of its known shortfalls is the assumption that future returns' moments are known with certainty, which has led to substantial under-performance in practice, as shown in [DGU09] and [CZ93]. Recently however there have been developments made to try to tackle this issue. Robust Optimisation is one of the techniques used to tackle this problem after having been successfully applied to other fields, see [FKPF07] and [BBC11].

Instead of specifying with certainty what the model parameters are, Robust Optimisation models only require a set of possible parameter values, the so-called uncertainty set. By far the most studied version of Robust Portfolio Optimisation is the worst-case scenario approach, whose aim is to find the portfolio with the best performance under the worst possible values in the uncertainty set. For a survey see [KKF14].

While this approach may be sensible in some situations, we believe it is not well suited for most practitioners. Even though it is important to worry about extreme scenarios, typical everyday scenarios are no less important and should not be ignored. Moreover, some professionals such as investment managers are frequently evaluated against the competition and not on absolute terms. For those reasons we believe that Relative Robust Optimisation, introduced in [KY97] and developed in [HKT13], is more suited to most portfolio managers. In this approach a portfolio's worth is assessed not only on its performance but also on the competition's performance.

Competition is clearly an ambiguous term which needs to be made precise. In [HKT13] the authors define competition as an "omniscient adversary" that has the same constraints as we do, but knows the "correct" model parameters and hence solves a non-robust portfolio optimisation problem. We introduce the idea that competition may be specified independently of our initial problem, i.e., we make it clear that competition may be described according to user's needs, just like the uncertainty set.

Most academic literature on Robust Portfolio Optimisation propose that the objective function must be replaced, which suggests investors must start anew if they wish to use such a tool. However investors have their own objectives and are wary of changing them, which makes the leap from literature to practice much harder to achieve. With this in mind we propose adding an extra constraint instead of a new objective function. This allows investors to keep their framework intact and simply make a minor change, trading an extra constraint for extra robustness. We also promote the use of *proportional regret*, a slightly different measure that is more in line with practitioners common practice.

In line with this philosophy we will look at models that deal with volatility, not expected returns. The idea of investing without taking into account expected returns is not new. Since [HB91] proposed a minimum volatility portfolio a lot of research has been published on it and its performance has been widely studied (see for example [CdST06] and [Lee11]). In fact, some asset managers already offer this solution to their clients in one form or another (see [Sch11]) and it was recently reported in the media that such funds had net inflows of $12.5 billion in the first half of 2016 (see [Kur16]). There has also been an increase in the literature of proposed metrics that solely look at risk, for example [TRM10] and [CC08].

It is however straightforward to extend our framework to models that incorporate expected returns or even higher moments. In fact this approach appears particularly promising given how sensitive traditional Portfolio Optimisation is to errors in the means, see [CZ93]. Covariance matrices in comparison are much more stable over time, although they do shift between different regimes, as shown in [FPW^{+}11]. We will assume a finite number of regimes, each characterised by a covariance matrix.

We start with a short review of relevant existing models: first the minimum volatility problem, followed by absolute robust portfolio optimisation and finally relative robust portfolio optimisation. We then move on to introduce our own model, by proposing a different measure of *regret* than the one found in [HKT13]. Initially we consider it an objective function, as is common in the literature, but subsequently we propose using it as a constraint instead, making it more widely applicable. We then pursue the suggestion in [KY97] and introduce *proportional regret* in the context of portfolio optimisation, a slightly different measure that is more in line with practitioners common practice. We briefly discuss how our model could be extended to a more general setting and at last we run the numerical experiments described above and we conclude with some final remarks.

It is worth mentioning that all proposed models can be efficiently solved in polynomial time with standard optimisation software. That is due to the fact that they can be cast as either second order cone programs or (for the more general version) semidefinite programs for which efficient algorithms are widely available.

2 Classical Models

We shall start with a quick look through some of the most common models found in the literature, before putting forward a different approach to deal with uncertainty in the covariance matrix.

Let $\mathcal{X} \subseteq \mathbb{R}^n$ denote the set of all admissible portfolios. This might simply be the set of all fully invested long-only equity portfolios, or it might be something different depending on individual constraints. [HB91] proposed one should invest in the minimum volatility portfolio, provided we know what the "true" $n \times n$ covariance matrix Q is. The problem of finding this *optimal* portfolio x can be written as

$$\begin{aligned} \min_{x} \quad & \sqrt{x^T Q\, x} \\ \text{s.t.} \quad & x \in \mathcal{X}. \end{aligned} \tag{1}$$

This model assumes Q is known for the coming investment period, which is not a reasonable assumption in financial applications. While there are many techniques available to provide good estimates, any approach of practical relevance is going to have to deal with the fact that market conditions change, as described in [FPW$^+$11], hence assuming a covariance matrix as certain is bound to fail. A natural step forward is to assume that the future covariance matrix is an unknown element of a set of *scenarios* $\mathcal{U} = \{Q_1, \ldots, Q_k\}$. This is the approach we shall take henceforth.

The challenge is how to make a decision when given such an uncertainty set. A possible avenue is to take a Bayesian approach by assigning each scenario with a probability p_i. This however requires us to come up a sensible vector of probabilities p, which will have a significant impact on the optimal portfolio x. We will therefore stay clear of this and similar approaches and seek instead an alternative that makes no distributional assumptions.

One such technique is *Absolute Robust Portfolio Optimisation*, which has been extensively studied in the literature, see for example [KKF14]. The aim is to minimise volatility under the worst-case scenario, i.e., under the scenario where volatility is highest for the chosen portfolio. The portfolio that satisfies this can be found by solving

$$\begin{aligned} \min_{x} \ \max_{Q \in \mathcal{U}} \quad & \sqrt{x^T Q\, x} \\ \text{s.t.} \quad & x \in \mathcal{X}. \end{aligned} \tag{2}$$

The problem is that this model focuses only on the worst-case scenario of each portfolio. While this might be attractive in subjects such as engineering, for financial markets the motivation is more lacking – people don't want to make investment decisions solely on the basis that everything will go wrong. This approach completely ignores all the other scenarios which means it will not be apt to reap rewards when a more positive scenario happens.

A more balanced approach is the one found in [HKT13], *Relative Robust Portfolio Optimisation*. The goal is to minimise the amount of extra volatility we are forced to accept for not knowing the covariance matrix with certainty. For each scenario $Q \in \mathcal{U}$, there is a portfolio y_Q which is optimal under that scenario (this can be found by solving (1)). Unfortunately we cannot know which scenario is correct beforehand, hence we do not know what the correct choice is. Any portfolio we choose will always lose out against the optimal portfolio of the true scenario. We can however strive to make it as low as possible, by minimising the worst outcome across all possible values of $Q \in \mathcal{U}$, that is, by solving

$$\begin{aligned} \min_{x} \max_{Q\in\mathcal{U}} \quad & \left(\sqrt{x^T Q\, x} - \min_{y\in\mathcal{X}} \sqrt{y^T Q\, y}\right) \\ \text{s.t.} \quad & x \in \mathcal{X}. \end{aligned} \tag{3}$$

All these methods build on the previous one and so shall we build on (3) by making the objective function less abstract and more practically focused. Our proposal shall allow a practitioner to deal with uncertainty without coming up with a probability distribution and without focusing solely on the worst-case scenario. At the same time it will measure a quantity of real interest, in this case the extra volatility we incur against certain benchmarks. All these traits set it apart from other methods in the literature, namely the ones previously mentioned.

3 Regret Minimisation

While (3) has potential, we wish to approach it differently. Comparing us against a "more knowledgeable version of ourselves" might be a very interesting theoretical problem, but in practice one gets compared against other players in the market, not fictional beings. Hence we cannot be satisfied by the formulation in (3).

We would like instead to be compared against other investors, as this would be a more realistic portrait of the real world. Unfortunately we do not know how other agents in the market are investing and therefore cannot hope to model precisely the behaviour of all investors.

As an alternative we will compare ourselves against benchmarks, which are naturally unambiguous. Benchmarks have two properties that make them suitable for this purpose. On the one hand, an asset manager is not only compared against better performing competitors, but also against relevant benchmarks. This means comparison against benchmarks is a reasonable representation of reality. On the other hand, competitors will themselves be victims of pressure not to fall too short of a particular benchmark. This will in turn pull them towards a similar basket, resulting in a high correlation between the two. Notice for example the historical correlation between the HFRI Equity Hedge Fund (Total) Index and the S&P 500 in Figure 1. This chart provides support to the claim that benchmarks are a good proxy for agents' returns. We can see that this particular basket of hedge funds have been increasingly correlated with the S&P 500, with recent values surpassing the 90% threshold.

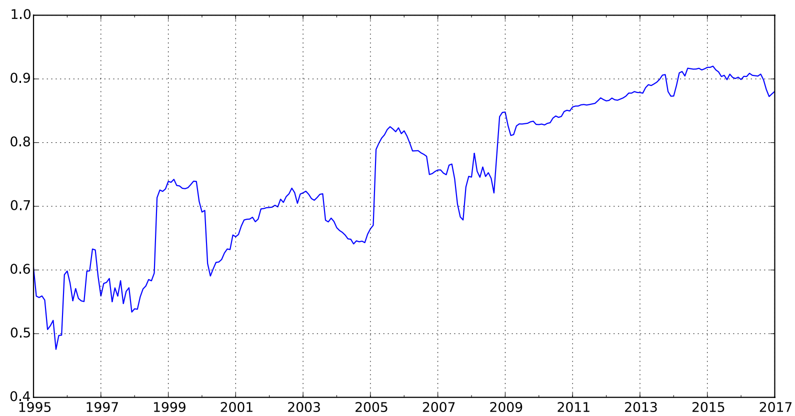

Figure 1: Correlation between the HFRI Equity Hedge (Total) Index and the S&P 500.

This is the case despite the fact that most hedge funds do not have explicit benchmarks – their mandate is to deliver positive returns and they are commonly considered a diversified investment.

3.1 Modelling Regret

Let $\mathcal{B} = \{b_1, \dots, b_m\} \subseteq \mathbb{R}^n$ be the set of benchmarks we wish to consider. Following the rationale behind (3), we wish to consider our excess volatility to the least volatile benchmark,

$$l_{\mathcal{B}}(x, Q) := \sqrt{x^T Q\, x} - \min_{b \in \mathcal{B}} \sqrt{b^T Q\, b}.$$

Here we assumed that all the underlyings of all benchmarks are included in the set of n assets considered for investment, and so we can compute their volatility by $\sqrt{b^T Q\, b}$ – in truth this is not strictly necessary, as well will see later on.

In a similar fashion to [HKT13] we define *regret* to be the maximum of this loss over $\mathcal{U}$, i.e.,

$$\operatorname{Rgrt}_{\mathcal{U},\mathcal{B}}(x) := \max_{Q \in \mathcal{U}} l_{\mathcal{B}}(x, Q) = \max_{Q \in \mathcal{U}} \left(\sqrt{x^T Q\, x} - \min_{b \in \mathcal{B}} \sqrt{b^T Q\, b} \right). \tag{4}$$

Regret can be understood as our distance to the "winner" under the least desirable scenario. In financial terms it measures how much more volatility we could be getting against the best performing benchmark – the bigger it is, the more we will have "regretted" not having invested similarly to that benchmark.

Benchmarks choice, just like scenarios, is beyond the scope of this paper. We can however provide a couple of useful tips that apply to both. First off, they should be picked independently of the portfolio optimisation problem, be it through economic/political considerations, data analysis, or any other technique. Furthermore, they should be as realistic and not too many. Too many scenarios and benchmarks will lead to overly defensive portfolios, since an "anything can happen" attitude is a sure way to investment paralysis.

If we wish to find the portfolio that has least regret then we have to solve the following *Relative Robust Portfolio Optimisation with Benchmarks* problem.

$$\begin{aligned} \min_{x} \quad & \operatorname{Rgrt}_{\mathcal{U},\mathcal{B}}(x) \\ \text{s.t.} \quad & x \in \mathcal{X}. \end{aligned} \tag{5}$$

Unlike the minimum volatility problem (1), minimum regret *might* be negative. It should be clear to the reader that this will be the case if and only if there exists a portfolio that has a lower volatility than all of the considered benchmarks regardless of which scenario is realised.

On a technical level, we actually do not need to have a set of defined benchmarks which we want to be compared against. All we need is a *target volatility* σ_i for each scenario Q_i. To show this we start by lifting problem

(5) via the introduction of an auxiliary variable γ,

$$\begin{aligned} \min_{x,\gamma} \quad & \gamma \\ \text{s.t.} \quad & \gamma \geq \max_{Q \in \mathcal{U}} \left(\sqrt{x^T Q\, x} - \min_{b \in \mathcal{B}} \sqrt{b^T Q\, b} \right). \\ & x \in \mathcal{X}. \end{aligned}$$

It is clear that $\gamma \geq \max_{Q \in \mathcal{U}} \left(\sqrt{x^T Q\, x} - \min_{b \in \mathcal{B}} \sqrt{b^T Q\, b} \right)$ is satisfied if and only if $\gamma \geq \sqrt{x^T Q\, x} - \min_{b \in \mathcal{B}} \sqrt{b^T Q\, b}$ for all $Q \in \mathcal{U}$. Moreover the quantity $\min_{b \in \mathcal{B}} \sqrt{b^T Q\, b}$ does not depend on x or γ and therefore can be calculated independently of the optimisation process. It is clear that it depends on Q, but since the set $\mathcal{U}$ is discrete we merely need to calculate that quantity for each single scenario. That quantity is our *target volatility*,

$$\sigma_i = \min_{b \in \mathcal{B}} \sqrt{b^T Q_i\, b}, \quad \forall i = 1, \ldots, k,$$

and we can hence write (5) as

$$\begin{aligned} \min_{x,\gamma} \quad & \gamma \\ \text{s.t.} \quad & \gamma \geq \sqrt{x^T Q_i\, x} - \sigma_i, \quad \forall i = 1, \ldots, k, \\ & x \in \mathcal{X}. \end{aligned}$$

Notice that there is no longer mention of the benchmark set $\mathcal{B}$ – all the relevant information was encoded in σ_i. Hence as an alternative to specifying $\mathcal{B}$ one can simply set a target volatility σ_i for each scenario Q_i.

While (5) can be expressed as a triple optimisation problem, it is actually a tractable problem that can be solved as efficiently as the minimum volatility problem (1), as long as $\mathcal{X}$ can be cast as a tractable cone (which is true for most admissible sets of practical relevance). This is made possible by the addition of auxiliary variables t (and γ) that allow (5) to be converted into

$$\begin{aligned} \min\ & \gamma \\ \text{s.t.}\ & \gamma - t_i + \sigma_i \geq 0, \quad \forall i = 1, \ldots, k, \\ & \begin{bmatrix} t_i \\ U_i^T x \end{bmatrix} \in L_{n+1}, \quad \forall i = 1, \ldots, k, \\ & x \in \mathcal{X}, \end{aligned} \tag{6}$$

where U_i is the Cholesky decomposition of $Q_i = U_i\, U_i^T$ and L_{n+1} is the Lorentz cone defined by

$$L_{n+1} = \left\{ \begin{bmatrix} t \\ x \end{bmatrix} \in \mathbb{R}^{n+1} : \|x\|_2 \leq t \right\}.$$

The optimisation problem found in (6) is a second order cone program and can thus be solved using standard solvers (such as CPLEX or MOSEK). This means that in the end we have a much richer approach that can be solved in roughly the same computational time as a standard Markowitz optimisation problem.

3.2 Regret as a Constraint

So far we have only seen models that assume volatility as the objective function. Another common approach is to use volatility instead as a constraint. In fact [Mar52] first introduced Modern Portfolio Theory as

$$\begin{aligned} \max_x \quad & \mu^T x \\ \text{s.t.} \quad & \sqrt{x^T Q\, x} \leq \kappa, \end{aligned}$$

where μ is the expected future returns and κ is the maximum volatility allowed in our portfolio. κ can thus be as low as the solution of (1).

Similarly, we can look at using regret not as an objective in itself, but rather as a constraint to be imposed in our portfolio allocation problem. Say for example we wish to maximise expected return as in the above problem. Then we could solve a similar problem,

$$\begin{aligned} \max_x \quad & \mu^T x \\ \text{s.t.} \quad & \text{Rgrt}_{\mathcal{U},\mathcal{B}}(x) \leq \theta, \end{aligned} \tag{7}$$

where θ is the regret we are prepared to accept and can be no lower than the solution to (5).

This approach is attractive on the grounds that investors prefer to target more meaningful objectives, such as maximising expected returns or minimising transaction costs. In addition it gives the portfolio manager control of how much risk he is prepared to take in order to increase (or decrease) his objective function. It is also a more flexible option, since adding a constraint is comparatively easier than rewriting from scratch an optimisation problem with a new objective function. In fact there is no reason why we cannot have both constraints – one on volatility and another on regret.

3.3 Proportional Regret

As an alternative to regret we could define excess volatility not as the difference in volatilities, but instead the proportion of benchmark volatility we surpass. We can then define *proportional regret* to be the maximum of this loss, i.e.,

$$\mathrm{PRgrt}_{\mathcal{U},\mathcal{B}}(x) := \max_{Q \in \mathcal{U}} \max_{b \in \mathcal{B}} \frac{\sqrt{x^T Q\, x} - \sqrt{b^T Q\, b}}{\sqrt{b^T Q\, b}}.$$

This was first sugested by [KY97] in the context of general robust optimisation. We believe thinking in proportions is more natural for investors and so this option will be of more practical interest to them. As before, the higher it is the more we "regret" not having invested similarly to the best performing benchmark. The difference here is that we do not measure the absolute increase in volatility but rather the proportional increase.

As before, we only require the volatility of the best benchmark under each scenario to be modelled, not for all of them. Likewise, proportional regret can only be negative if x outperforms all benchmarks over all scenarios.

It should be made clear, however, that this is not just a cosmetic change. While the best benchmark in each scenario remains the same, the trade-off between scenarios does change. To illustrate this let us look at a quick example.

Imagine we have two assets we can invest in, but we can only invest in one of them. We will be compared against a single benchmark – an equally weighted basket of the two assets. Imagine further that over the next year there are two possible scenarios, summarised in the following covariance matrices

$$Q_1 = \begin{bmatrix} 0.02 & 0.006 \\ 0.006 & 0.05 \end{bmatrix}, \qquad Q_2 = \begin{bmatrix} 0.07 & 0 \\ 0 & 0.05 \end{bmatrix}.$$

The first asset has $\sqrt{0.02} \approx 14\text{pp}$[1] volatility in Scenario 1 and 26pp volatility in Scenario 2, while the second asset has 22pp volatility in both Scenario 1 and 2. In addition we have a single benchmark b that has 14pp volatility in Scenario 1 and 17pp volatility in Scenario 2.

One can check that Asset 1 performs worst in Scenario 2, while Asset 2 performs worst in Scenario 1. In Scenario 2 Asset 1 trails b by 9pp, or 53%, while in Scenario 1 Asset 2 trails b by 8pp, or 57%. It should then be

[1]To avoid confusion between Regret and Proportional Regret we will refer to volatility units as percentage points (pp).

clear that if we cared about regret we would invest in Asset 2, while if we cared about proportional regret we would invest in Asset 1. Hence these two approaches to measuring distance to benchmarks yield different decisions.

Proportional regret can be used as an objective simply by substituting $\mathrm{Rgrt}_{\mathcal{U}}$ by $\mathrm{PRgrt}_{\mathcal{U}}$ in (5). It can also be used as a constraint, in which case we make the same substitution in (7) instead.

Just like (5), all these variations can be efficiently solved using standard tools. For example, a maximum expected return problem with a proportion regret constraint can be shaped into the following second order cone program, where θ is once again a risk budget that has to be chosen large enough for the programme to have feasible points.

$$\begin{aligned}
\min\ & \mu^T x \\
\text{s.t.}\ & (1+\theta)\,\sigma_i - t_i \geq 0, \quad \forall i = 1, \ldots, k, \\
& \begin{bmatrix} t_i \\ U_i^T x \end{bmatrix} \in L_{n+1}, \quad \forall i = 1, \ldots, k, \\
& x \in \mathcal{X}.
\end{aligned}$$

4 General relative robust optimisation problem with benchmarks

The idea of comparing ourselves against benchmarks can be extended to any relative robust optimisation problem. Let $f(x;\lambda)$ be a parametric objective function we wish to minimise, where λ is an unknown parameter vector we have no control over, but know (or believe) to be in $\mathcal{U}$. Then we can define regret as

$$\mathrm{Rgrt}_{\mathcal{U},\mathcal{B}}\,(x) = \max_{\lambda \in \mathcal{U}} \max_{b \in \mathcal{B}} f(x;\lambda) - f(b;\lambda),$$

where $\mathcal{B}$ is the set of benchmarks we wish to consider. If f is strictly positive over $\mathcal{B}$ we can alternatively consider proportional regret, as defined by

$$\mathrm{PRgrt}_{\mathcal{U},\mathcal{B}}\,(x) = \max_{\lambda \in \mathcal{U}} \max_{b \in \mathcal{B}} \frac{f(x;\lambda) - f(b;\lambda)}{f(b;\lambda)}.$$

Formally the difference to the definitions found in [KY97] is the substitution of the permissible values for x, the admissible set $\mathcal{X}$, for a general set $\mathcal{B}$. In practice, however, the differences in motivation are important – classical relative robustness compares us against *what we could have done had we guessed the correct parameters and optimised accordingly.* We propose that we merely compare ourselves against real benchmarks. This has special importance in financial applications, where direct or indirect comparison to established benchmarks is widespread.

For the sake of simplicity we will restrict our analysis to the minimum regret optimisation problem (5). Similar conclusions can be taken for the other proposed models. We can write this problem as

$$\begin{aligned} &\min \gamma \\ &\text{s.t. } \gamma - f(x;\lambda) + f(b;\lambda) \geq 0, \quad \forall \lambda \in \mathcal{U},\ b \in \mathcal{B}, \\ &\qquad x \in \mathcal{X}. \end{aligned} \tag{8}$$

One thing worth pointing out is that there is no reason as to why $\mathcal{U}$ and $\mathcal{B}$ have to be independent sets – this means that we can just as easily pick different benchmarks for different scenarios. From now on we will consider the set Θ as the set of doubles (λ, b) such that b is a benchmark we wish to consider if the true parameter is λ. We can then substitute (8) for

$$\begin{aligned} &\min \gamma \\ &\text{s.t. } \gamma - f(x;\lambda) + f(b;\lambda) \geq 0, \quad \forall (\lambda, b) \in \Theta, \\ &\qquad x \in \mathcal{X}. \end{aligned} \tag{9}$$

It is worth pointing out that from Θ we can construct sets $\mathcal{U}$ and $\mathcal{B}$ by projecting Θ into its first and second component, respectively. However in general solving (9) will not be the same as solving (8), as long as $\Theta \neq \mathcal{U} \times \mathcal{B}$.

If $\mathcal{U}$ and $\mathcal{B}$ are discrete sets, then the first constraint can be immediately broken down into a finite number of simpler constraints and the problem can be solved efficiently for most sensible functions f. The question is when can we efficiently solve (9) when $\mathcal{U}$ and $\mathcal{B}$ are closed, convex sets. To answer that question we will make the extra assumption that f is a concave quadratic function in (x, λ) but linear in x. This might sound too restrictive, but usually this is made possible by the introduction of auxiliary variables, as was done in (6).

Under these conditions, the first constraint can be written as

$$\gamma - f(x;\lambda) + f(b;\lambda) := q_{x,\gamma}(\lambda, b) \in \mathscr{FC}_+(\Theta),$$

where $\mathscr{FC}_+(\Theta)$ is the set of quadratics (or their matrix representations) that are positive over Θ. Since $q_{x,\gamma}$ is linear in (x, γ), the only question that remains is whether $\mathscr{FC}_+(\Theta)$ is a tractable set. Unfortunately this is not the case in general – in fact for some Θ the problem of deciding whether a matrix A is in $\mathscr{FC}_+(\Theta)$ is co-NP-complete, see [MK87].

There is however some hope. If Θ is the intersection of n_e, ellipsoids, n_p hyperplanes and n_h half-planes, then one can efficiently solve (9) if n_e and n_p are at most 1. If either (or both) $n_e \geq 2$ or $n_p \geq 2$ then one has to substitute $\mathscr{FC}_+(\Theta)$ by an inner approximation, which in turn means we can only get an upper bound on (9). For $n_e = 1$ such an inner approximation is derived in [HT07], but a similar derivation works for $n_e > 1$.

While inner approximations might not feel fully satisfactory, it is possible to attain not only the upper bound efficiently but also a portfolio that guarantees it – so this model can still be used in more general contexts.

5 Empirical Results

In this section we will investigate if the proposed method yields sensible results and, equally importantly, if the results are fundamentally different from other approaches commonly used in the literature. All these models can be solved efficiently using a standard conic optimisation solver, such as MOSEK or Gurobi.

To do this we will investigate its use in building a fully-invested long-only portfolio of equities across different regions. Since our goal is mostly illustrative we choose to regard as investable asset industry group indices instead of individual companies. For that purpose we follow the recommendations of [BLO03] and use the Global Industry Classification Standard (GICS). We will therefore consider each of the 24 industry groups in GICS as investable assets.

Benchmark choice is a delicate process that will vary substantially from one investor to the next. Given this, we will take, for simplicity's sake, advantage of the hierarchical structure of GICS and use as benchmarks its 10 sectors along with the region wide total return index. We will use data from 2001 onwards, with prices quoted in local currency, and for each specific region will only consider industry groups and sectors that had a non-interrupted presence.

Another choice that needs to be made is what are the scenarios we wish to consider. There exist an almost countless number of ways to generate scenarios that take into account all kinds of data and/or expertise. We made the decision to solely use historical data up to 2012 in order to generate our scenarios, which we will then use to compute portfolios under different portfolio optimisation techniques. Data from 2013 to 2016 will then be used to analyse these portfolios' performance. The scenarios will be covariance matrices of subperiods of our training horizon. To determine the changepoints we implement the algorithm proposed in [BO16]. This is an attractive option since it makes no distribution assumption and so fits well in the spirit of these experiments. While we could in theory have “forced” some dates to be changepoints (Lehman Brothers collapse or 9/11, just to name a couple of possibilities), we wish to have as little direct influence as possible and so will stick with the dates generated by the algorithm.

5.1 Minimum Regret in the EMU

We will start by focusing on the European Economic and Monetary Union equity market. After running the algorithm proposed in [BO16], we arrive at the following changepoints: 31/5/02, 06/06/03, 22/11/04, 27/04/06, 26/11/08, 26/04/10 and 04/01/12. From these we construct 8 sub-periods, which we then summarise into 8 covariance matrices to be our scenarios. We can visualise them through the correlation heatmaps in Figure 2.

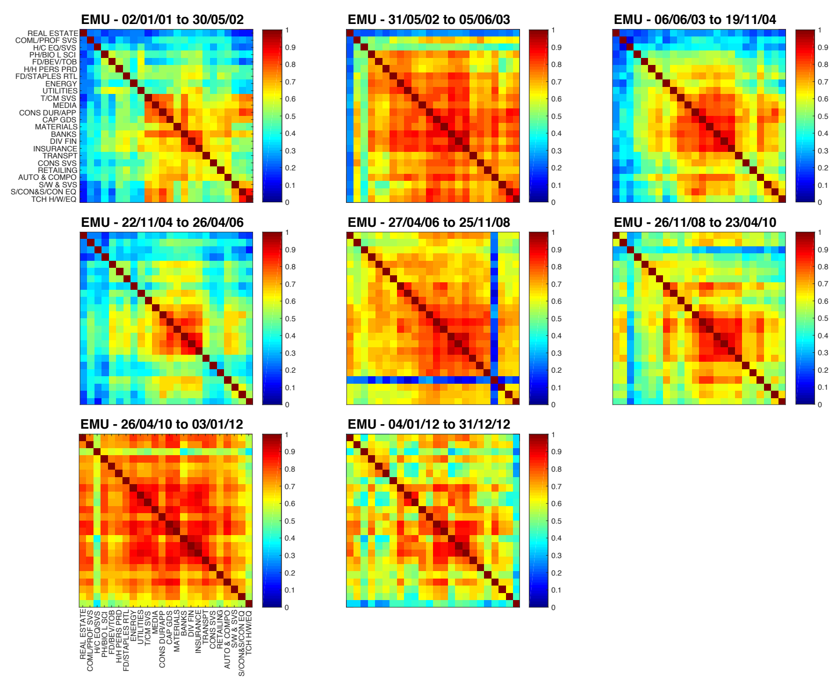

Figure 2: Heatmaps of correlation matrix for different scenarios in the EMU.

One can check whether these scenarios are sensible by following the methodology proposed in [FPW+11]. In Figure 3 we can see the behaviour of the first principal component of a one-year rolling window of returns over different scenarios. The first principal component is the market direction that accounts for the most variance, hence looking at how much of the total variance in the market is explained by a single direction is a powerful tool to visualise how fragmented the market is.

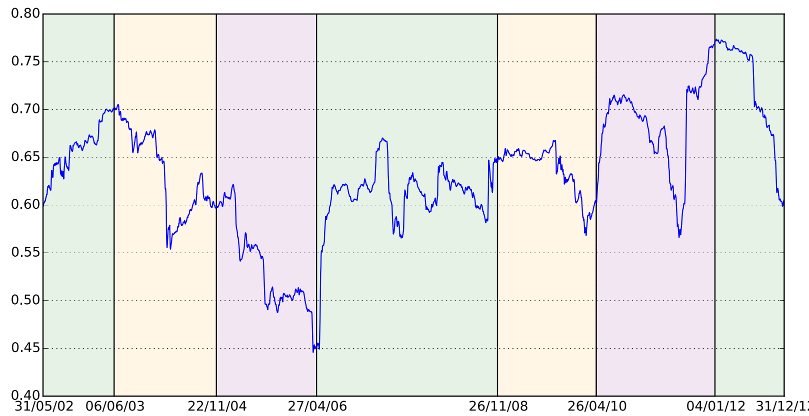

Figure 3: Proportion of variance explained by the first principal component in EMU over different scenarios.

We can see that the changepoint detection algorithm is picking up genuine changes in the first principal component, hence reinforcing our conviction that these constitute true regime changes. The first scenario is purposely absent as the use of a one-year rolling window rules out this analysis over the first year, which is over two-thirds of the total time spent in that scenario.

We are now in a position to run our relative robust portfolio optimisation (RRPO), as in (5), and compare it to other methods. Two methods that are reasonable to compare our model against are the absolute robust portfolio optimisation (ARPO), as in (2), using the same scenarios as RRPO, and a simple minimum volatility portfolio optimisation (MVPO), as in (1), using the whole training period to compute the covariance matrix estimate. These portfolios have their weights detailed in Table 1.

	RRPO	ARPO	MVPO
Food, Beverage & Tobacco	35.84%	32.96%	36.77%
Health Care Equipment & Services	33.28%	50.23%	44.17%
Food & Staples Retailing	10.15%	0%	0%
Media	10.13%	0%	0%
Telecommunication Services	5.74%	0%	2.16%
Real Estate	4.86%	4.09%	12.26%
Transportation	0%	9.65%	4.64%
Automobiles & Components	0%	3.07%	0%

Table 1: Weights of different portfolios being considered for EMU.

Upon examining this table, one can immediately draw a few conclusions. First, all three portfolios concentrate on a small subset of all investable options (GICS has 24 industry groups, whereas Table 1 makes mention of only 8 of them). This is hardly surprising, as all of them are aiming at the same objective: having low volatility. It is then only natural that they tend to concentrate on low volatility assets. Second, RRPO is slightly more diversified than ARPO and MVPO, with a maximum weight of 35.84% over 6 assets against a maximum of 50.23% and 44.17% over 5 of the other methods. Finally, it is evident that RRPO is yielding a structurally different portfolio to both ARPO and MVPO – the proportion of capital allocated differently is 29.67% and 23.86% respectively, both non-trivial quantities.

Let us now have a look at how these portfolios perform out of sample, over the period starting from 2013 until 2016. In Figure 4 we have a plot of the total return over time for each portfolio, while in Table 2 we list the realised return and volatility, in yearly terms, of each portfolio.

It so turns out that for this dataset all three methods behave remarkably similarly, even if one takes into account the fact that roughly 75% of capital is allocated equally. As we will see later on, this behaviour is purely coincidental – the same methods across different regions have discernible differences in realised return. We can however take note that RRPO was the least volatile portfolio on any of the years considered.

5.2 Maximum Expected Return in the US

We now wish to test another of the proposed methods, a proportional regret constraint of 10% in a maximum expected return portfolio optimisation

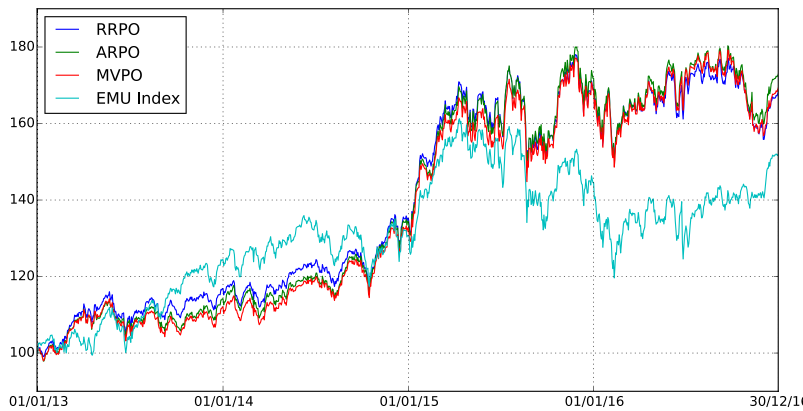

Figure 4: Performance of RRPO, ARPO, MVPO and EMU Index.

(PRCPO), as in (7) but with $\text{PRgrt}_{\mathcal{U}}$ instead of $\text{Rgrt}_{\mathcal{U}}$. This will give us an example of a situation where the objective is already set, and we solely add a constraint to force a more stable behaviour. We will compare this against a Sharpe ratio absolute robust portfolio optimisation (SARPO) (where we aim to have the best worst-case Sharpe ratio) and against a maximum Sharpe ratio portfolio optimisation (MSPO). Moreover we will see how these methods compare when there are other restrictions present. We will not only keep the long-only constraint, but we will also add a 20% cap to force diversification.

We will use US equity market data so we can look in detail at another dataset. Our investable assets will still be the 24 GICS industry groups and our benchmarks the 10 GICS sector groups plus the US Index. The training and testing period also remain unchanged.

As before, we use the changepoint detection proposed in [BO16] to come up with our scenarios. The changepoints detected are 31/05/02, 15/06/04, 26/02/07, 03/09/08 and 20/12/11. We can see the heatmaps of each individual scenario in Figure 5.

As before, we can see the behaviour of the first principal component

	RRPO	ARPO	MVPO	EMU Index
Realised Return 2013	15.04%	12.80%	10.97%	21.09%
Realised Return 2014	14.13%	15.79%	16.34%	4.78%
Realised Return 2015	21.45%	23.66%	23.19%	9.70%
Realised Return 2016	-0.34%	0.65%	0.40%	4.90%
Realised Volatility 2013	11.95pp	12.33pp	12.23pp	14.29pp
Realised Volatility 2014	12.36pp	13.10pp	12.72pp	15.41pp
Realised Volatility 2015	20.94pp	21.55pp	21.45pp	21.21pp
Realised Volatility 2016	17.46pp	18.42pp	17.98pp	20.30pp

Table 2: Total return series assuming initial capital of 100. Realised return and volatility calculated for each civil year.

over different scenarios, except the first, in Figure 6. We can once again see clear differences in behaviour, which justifies our choice of scenarios as true regime changes. It is worth pointing out that while Figure 6 might suggest we should either anticipate or delay some of the changepoints, their calculation was made using the whole covariance information, not just the first principal component. Making decisions based on Figure 6 would throw away all the information contained in the remaining principal components and would hence be suboptimal.

We use the 6 covariance matrices for both PRCPO and SARPO. For MSPO we use the covariance matrix calculated over the whole training period as our estimate. For all three portfolio optimisation problems we will use as the expected return of each individual asset the realised return over the most recent scenario – in this case, the realised return between 20/12/11 and 31/12/12. This is an arbitrary choice that is by no means a superior predictor to any other prediction, but is simple enough as a rule of thumb for us to use it as an example. Not only is forecasting expected returns well beyond the scope of this paper, but they also have a disproportionate impact on portfolio weights when compared to expected volatility or correlation, as demonstrated in [DGU09] and [Mic89]. Hence different forecasting techniques are likely to considerably impact the results that follow. We therefore urge readers to go the extra mile if they are committed to incorporating expected returns in portfolio optimisation problems.

After solving these 3 optimisation problems, the resulting portfolios can be found in Table 3.

It is unequivocal that different choices have been made, despite the fact that the objective function (i.e. the expected returns) is the same for all

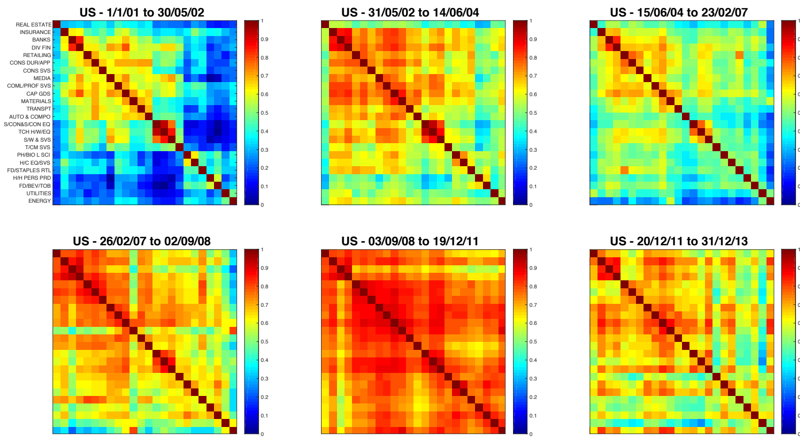

Figure 5: Heatmaps of correlation matrix for different scenarios in the US.

three portfolios. PRCPO focuses more on Food, Beverage & Tobacco, while SARPO and MSPO focus more on Telecommunication Services, Media and Retailing. They are all fairly diversified, but this is merely a symptom of the imposed cap. PRCPO agrees with both SARPO and MSPO on 60.56%, meaning 39.44% of the capital is allocated differently.

We can now analyse how these portfolios perform on the test period in Figure 7 and Table 4.

We can see that SARPO has a slightly higher annualised Sharpe ratio, but this however comes at the cost of higher volatility than we would consider admissible (this is relevant as we require the portfolio to be fully invested). In 2014, for example, SARPO and MSPO had respectively 17.2% and 13.4% more volatility than the least volatile benchmark, the Consumer Staples sector group, while PRCPO had only 1.93% extra volatility, well below the 10% cap imposed. While both SARPO and MSPO did outperform the market over this period, it could easily have gone the other way. This is a good example of why we think a proportional regret constraint is important – in this case it allowed us to pursue a high Sharpe Ratio while safeguarding

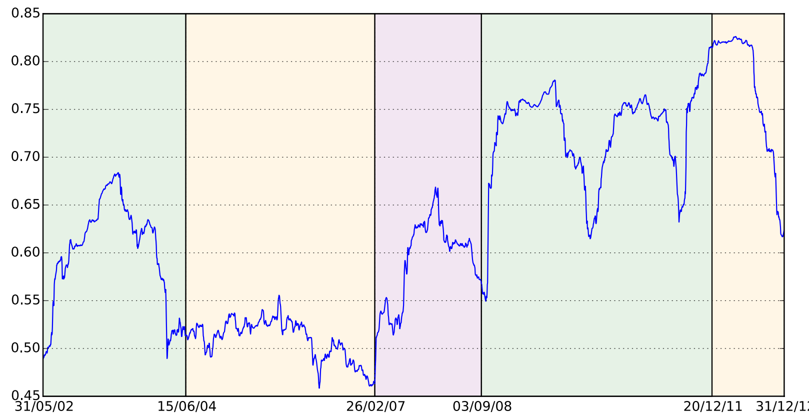

Figure 6: Proportion of variance explained by the first principal component in US over different scenarios.

against taking too much extra volatility.

5.3 Model Performance over Multiple Regions

As a way of reinforcing the points made earlier we carried out the same experiments on equity data from Australia, Canada, France, Japan, South Korea and United Kingdom. For each we calculated new changepoints and recomputed the portfolios above.

Tables 5 and 6 in the Appendix summarise the results for each region. In the minimum regret problem we outline the percentage of capital that is allocated differently from RRPO to ARPO and MVPO, the biggest weight each allocates to a single industry group, their realised return and volatility for the entirety of the testing period. In the maximum expected value problem we display the percentage of capital allocated differently from PRCPO to SARPO and MSPO, their Sharpe ratio and volatility (the maximum weight is 20% by construction).

	PRCPO	SARPO	MSPO
Food & Staples Retailing	20%	20%	20%
Food, Beverage & Tobacco	20%	0%	0%
Pharmaceuticals, Biotechnology & Life Sciences	20%	20%	20%
Household & Personal Products	19.44%	0%	0%
Telecommunication Services	12.94%	20%	20%
Media	7.62%	20%	20%
Retailing	0%	19.12%	9.40%
Real Estate	0%	0%	5.93%
Diversified Financials	0%	0%	4.66%
Health Care Equipment & Services	0%	0.88%	0%

Table 3: Weights of different portfolios being considered for US.

A few important points can be taken from this table. First, it is clear that both RRPO and PRCPO generally offer different solutions not covered by other methods. Secondly, RRPO tends to offer more diversification (by having less weight in a single industry group) than both ARPO and MVPO. In fact, the only instances when it does not is when the amount of capital allocated differently is marginal. This comes naturally from the fact that RRPO takes all scenarios into account, not only the worst-case scenario (as ARPO) and so favours diversification.

Thirdly, the only time one strategy significantly outperforms the others is PRCPO outperforming both SARPO and MSPO in Canada. Finally, it is worth mentioning that over the past few years low volatility assets have outperformed the market. This means that portfolios that invest in these assets, such as the ones examined, have also outperformed the market. There is no reason to believe this to be indefinitely true, and so none of these portfolios is guaranteed to persistently outperform the market.

6 Final Remarks

In this paper we introduced a variation of Relative Robust Portfolio Optimisation that we believe can have a significant impact in the current use of portfolio optimisation by practitioners. By considering a defined set of benchmarks we give a tangible intuition which was previously lacking to the notion of regret, making clearer the case that there are better options in

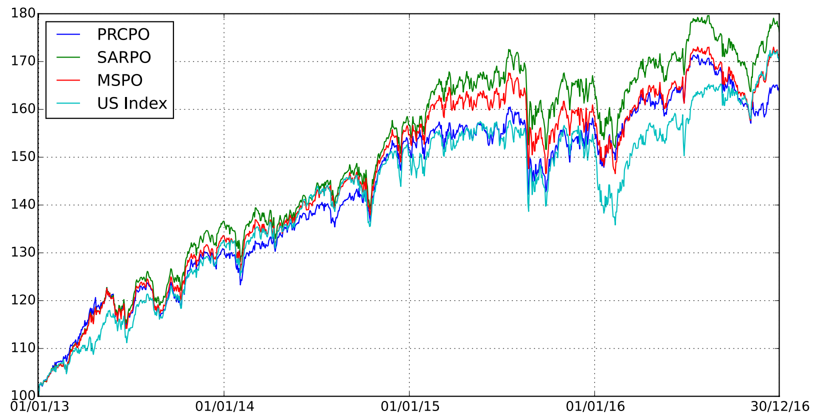

Figure 7: Performance of PRCPO, SARPO, MSPO and US Index.

Robust Portfolio Optimisation than a worst-case scenario approach.

We also propose other innovations that aim to address common concerns in the industry. A natural risk aversion to change amongst practitioners makes the adoption of brand new techniques very difficult, especially when they require changes in tried and tested systems. Being able to use regret as a constraint makes adoption much more easier, as it can added or removed from an already existing portfolio optimisation problem with ease. Similarly the use of measures practitioners are more familiar with and understand better makes our model less foreign and more likely to be well received.

Theoretically this is also a very interesting problem, as it is not yet clear which variants can be efficiently solved and which can merely be approximated. At the moment when taking $\mathcal{B}$ and $\mathcal{U}$ to be convex sets we are limited to one hyperplane and one ellipsoid, else an approximation is required. It is however not clear if it is impossible to efficiently solve for any higher number or if on the contrary such algorithms do exist. More research is thus needed on this topic.

We also looked at how two potential strategies fared against other common

	PRCPO	SARPO	MSPO	US Index
Realised Sharpe Ratio 2013	2.40	2.78	2.56	2.50
Realised Sharpe Ratio 2014	1.48	1.20	1.32	1.08
Realised Sharpe Ratio 2015	0.20	0.33	0.17	0.08
Realised Sharpe Ratio 2016	0.44	0.57	0.61	0.82
Realised Volatility 2013	10.80pp	10.82pp	10.94pp	10.88pp
Realised Volatility 2014	9.61pp	11.05pp	10.69pp	11.22pp
Realised Volatility 2015	13.69pp	14.52pp	14.35pp	15.15pp
Realised Volatility 2016	10.46pp	11.52pp	11.48pp	12.97pp

Table 4: Total return series assuming initial capital of 100. Realised return and volatility calculated for each civil year.

approaches in the literature over a range of equity index markets. Regret minimisation seems to provide a greater degree of protection across different regions when compared to absolute robust optimisation and the more widely used minimum volatility optimisation. It does assume a clean break from any previous investment strategy, so it should be of primary interest to academics, although it may also be implemented by a defensive portfolio manager whose main concern is not to be too heavily outperformed. A proportional regret constraint constructed using a suitable set of benchmarks and scenarios, however, is likely to be better suited towards investors, as it gives them the confidence to pursue their objectives while at the same time controlling the level of extra volatility they are prepared to take. Its numerical experiments are somewhat inconclusive because maximum return optimisation problems such as (7) are much more dominated by the expected return vector μ than by any other single feature. Nonetheless we think it provides a worthy risk management option for any portfolio manager, especially one who is measure against well known benchmarks.

References

[BBC11] D. Bertsimas, David B. Brown, and Constantine Caramanis. Theory and applications of robust optimization. *SIAM Review*, 53(3):464–501, 2011.

[BLO03] Sanjeev Bhojraj, Charles M. C. Lee, and Derek K. Oler. What's my line? a comparison of industry classification schemes for capital market research. *Journal of Accounting Research*, 41(5):745–774, 2003.

[BO16] Ian Barnett and Jukka-Pekka Onnela. Change point detection in correlation networks. *Scientific reports*, 6, (2016), 2016.

[CC08] Yves Choueifaty and Yves Coignard. Toward maximum diversification. *The Journal of Portfolio Management*, 35(1):40–51, 2008.

[CdST06] Roger G. Clarke, Harindra de Silva, and Steven Thorley. Minimum- variance portfolios in the u.s. equity market. *The Journal of Portfolio Management*, 33(1):10–24, 2006.

[CZ93] Vijay Kumar. Chopra and William T. Ziemba. The effect of errors in means, variances, and covariances on optimal portfolio choice. *The Journal of Portfolio Management*, 19(2):6–11, 1993.

[DGU09] Victor DeMiguel, Lorenzo Garlappi, and Raman Uppal. Optimal versus naive diversification: How inefficient is the 1/n portfolio strategy? *The Review of Financial Studies*, 22(5):1915–1953, 2009.

[FKPF07] Frank J. Fabozzi, Petter N. Kolm, Dessislava A. Pachamanova, and Sergio M. Focardi. Robust portfolio optimization. *The Journal of Portfolio Management*, 33(3):40–48, 2007.

[FPW+11] Daniel J. Fenn, Mason A. Porter, Stacy Williams, Mark McDonald, Neil F. Johnson, and Nick S. Jones. Temporal evolution of financial-market correlations. *Physical Review E: Statistical, Nonlinear, Biological, and Soft Matter Physics*, 84(2), 2011.

[HB91] Robert A. Haugen and Nardin L. Baker. The efficient market inefficiency of capitalization-weighted stock portfolios. *The Journal of Portfolio Management*, 17(3):35–40, 1991.

[HKT13] R. Hauser, V. Krishnamurthy, and R. H. Tütüncü. Relative robust portfolio optimization. *arXiv, (2013)*, 2013.

[HT07] R. Hauser and R. Tütüncü. Sdp approximability of relative robust qp. *PAMM*, 7(1):2060043–2060044, 2007.

[KKF14] Jang Kim, Woo Kim, and Frank Fabozzi. Recent developments in robust portfolios with a worst-case approach. *Journal of Optimization Theory and Applications*, 161(1):103–121, 2014.

[Kur16] Aaron Kuriloff. Stock investors pay up for peace of mind. *The Wall Street Journal*, 2016.

[KY97] Panos Kouvelis and Gang Yu. *Robust Discrete Optimization and Its Applications*. Kluwer Academic Publishers, Norwell, MA, 1997.

[Lee11] Wai Lee. Risk- based asset allocation: A new answer to an old question? *The Journal of Portfolio Management*, 37(4):11–28, 2011.

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

[Mic89] R. O. Michaud. The markowitz optimization enigma: Is 'optimized' optimal? *Financial Analysts Journal*, 45(1):31–42, 1989.

[MK87] Katta Murty and Santosh Kabadi. Some np- complete problems in quadratic and nonlinear programming. *Mathematical Programming*, 39(2):117–129, 1987.

[Sch11] Bernd Scherer. A note on the returns from minimum variance investing. *Journal of Empirical Finance*, 18(4):652–660, 2011.

[TRM10] Jé Teiletche, Thierry Roncalli, and Sé Maillard. The properties of equally- weighted risk contributions portfolios. *IDEAS Working Paper Series from RePEc*, 2010.

Appendix: Results across different regions

	AUS	CAN	FRA	JPN	KOR	UK	EMU	US
A1	30.51%	2.33%	33.16%	33.59%	36.32%	5.85%	29.67%	12.62%
A2	20.26%	14.36%	28.85%	14.45%	8.97%	3.07%	23.86%	9.54%
B1	25.15%	38.96%	19.22%	25.12%	35.49%	32.37%	35.84%	53.22%
B2	29.04%	39.74%	26.38%	43.06%	43.60%	33.07%	50.23%	58.37%
B3	28.24%	41.83%	25.84%	30.13%	39.26%	31.78%	44.17%	49.31%
C1	9.64%	12.48%	8.06%	16.15%	4.14%	7.77%	12.57%	11.82%
C2	11.82%	12.44%	6.75%	16.59%	4.13%	8.01%	13.23%	12.22%
C3	10.99%	14.22%	6.83%	15.72%	4.36%	7.46%	12.73%	11.96%
C4	9.11%	8.21%	10.60%	15.07%	1.73%	7.94%	10.12%	12.81%
D1	12.73pp	11.14pp	16.82pp	19.55pp	12.70pp	13.09pp	16.12pp	11.43pp
D2	13.08pp	11.23pp	16.52pp	19.30pp	14.23pp	13.09pp	16.79pp	11.41pp
D3	12.90pp	10.93pp	16.44pp	19.46pp	12.86pp	13.14pp	16.55pp	11.31pp
D4	14.07pp	12.01pp	18.23pp	22.22pp	12.99pp	14.34pp	18.06pp	12.68pp

Table 5: Behavior of RRPO across 8 different regions, from 2013 to 2016:
A - Percentage of capital RRPO invests differently from ARPO (1) and MVPO (2),
B - Highest weight in RRPO (1), ARPO (2) and MVPO (3),
C - Realised return of RRPO (1), ARPO (2), MVPO (3) and Regional Index (4),
D - Realised volatility of RRPO (1), ARPO (2), MVPO (3) and Regional Index (4).

	AUS	CAN	FRA	JPN	KOR	UK	EMU	US
E1	23.21%	35.58%	24.41%	22.24%	0.00%	15.32%	18.49%	39.44%
E2	40.00%	43.57%	11.74%	19.00%	8.57%	15.54%	29.51%	39.44%
F1	0.90	0.94	0.58	0.80	0.32	0.79	0.74	1.06
F2	0.99	0.37	0.51	0.72	0.32	0.70	0.76	1.13
F3	0.96	0.53	0.56	0.78	0.36	0.79	0.69	1.08
F4	0.65	0.68	0.58	0.68	0.13	0.55	0.56	1.01
G1	13.24pp	10.96pp	17.15pp	19.75pp	13.89pp	13.73pp	16.70pp	11.26pp
G2	12.65pp	17.38pp	17.37pp	20.63pp	13.89pp	14.36pp	16.72pp	12.08pp
G3	12.38pp	17.35pp	16.61pp	19.70pp	13.13pp	14.73pp	16.49pp	11.97pp
G4	14.07pp	12.01pp	18.23pp	22.22pp	12.99pp	14.34pp	18.06pp	12.68pp

Table 6: Behavior of PRCPO across 8 different regions, from 2013 to 2016:
E - Percentage of capital PRCPO invests differently from SARPO (1) and MSPO (2),
F - Realised Sharpe ratio of PRCPO (1), SARPO (2), MSPO (3) and Index (4),
G - Realised volatility of PRCPO (1), SARPO (2), MSPO (3) and Regional Index (4).