

Stability-based, random matrix theory filtering of financial portfolios

Justin Daly

Bachelor of Arts in Mathematics

A Dissertation submitted in fulfillment of the
requirements for the award of
Doctor of Philosophy (Ph.D.)

to the



Dublin City University

Faculty of Engineering and Computing, School of Computing

Supervisors: Prof. Heather J. Ruskin, Dr. Martin Crane.

September 2009

Declaration

I hereby certify that this material, which I now submit for assessment on the programme of study leading to the award of Doctor of Philosophy is entirely my own work, that I have exercised reasonable care to ensure that the work is original, and does not to the best of my knowledge breach any law of copyright, and has not been taken from the work of others save and to the extent that such work has been cited and acknowledged within the text of my work.

Signed

Justin Daly

Student ID 55131590

Date

Contents

Abstract	viii
Acknowledgements	ix
List of Tables	x
List of Figures	xv
1 Introduction	1
1.1 Motivation	1
1.2 Objectives	2
1.3 Thesis Outline	4
2 Literature Review	6
2.1 Introduction	6
2.2 Portfolio Theory and Noise	8
2.3 Volatility Forecasting and Riskmetrics	9
2.4 Applications of Some Aspects of Random Matrix Theory to Portfolio Se- lection	11
2.4.1 Work of Laloux et al. (1999; 2000)	11
2.4.2 Work of Plerou et al. (1999; 2000a; 2000b; 2001; 2002)	14
2.4.3 Agreement between Laloux et al. (1999; 2000) and Plerou et al. (1999; 2000a; 2000b; 2001; 2002)	15

2.4.4	Other Work on RMT for Large Stock Portfolios	16
2.4.5	Work of Pafka and Kondor (2002a,b) and Pafka et al. (2004)	17
2.4.5.1	Pafka and Kondor (2002a,b)	17
2.4.5.2	Pafka et al. (2004)	18
2.4.6	Small Portfolios	20
2.5	Stability Filtering	21
2.6	Eigenvalue Analysis	23
2.7	Summary	24
3	Background and Methodologies	25
3.1	Introduction	25
3.2	Background	25
3.2.1	Hedge Fund Portfolio Investment	25
3.2.2	Foreign Exchange Market	26
3.2.3	RMT and Historical Covariance	27
3.2.4	RMT and Exponentially Weighted Covariance	28
3.2.5	Standard Filtering Methods	30
3.2.5.1	LCPB Filtering Method	31
3.2.5.2	PG+ Filtering Method	32
3.2.6	Stability Based Filtering	32
3.2.6.1	Krzanowski Stability	32
3.2.6.2	Stability-based KR Filter	34
3.3	Methodologies	36
3.3.1	In-sample Methodology	36
3.3.2	Out-of-sample Methodology	40
3.4	Summary of Scenarios Tested	43
3.5	Summary	44

4	RMT for Large Stock Portfolios	46
4.1	Introduction	46
4.2	Review of Pafka et al.	47
4.3	Data	48
4.4	In-sample Testing	49
4.4.1	Methodology	49
4.4.2	Measuring Noise	49
4.4.3	Measuring Stability	52
4.4.4	In-sample Risk Reduction	53
4.5	Out-of-sample Testing	59
4.5.1	Methodology	59
4.5.2	Out-of-sample Analysis	59
4.5.2.1	Overall Risk	59
4.5.2.2	Annual Risk	59
4.5.2.3	Daily Risk	61
4.5.3	Parameter Values	62
4.5.4	Range of Realised Risks	64
4.6	Filtering Covariance	68
4.6.1	In-sample Analysis	68
4.6.2	Out-of-sample Analysis	73
4.6.2.1	Overall Risk	73
4.6.2.2	Annual Risk	74
4.6.2.3	Daily Risk	74
4.6.3	Parameter Values	76
4.6.4	Range of Realised Risks	77
4.7	Comparison of Tested Filters	80
4.7.1	In-sample Comparison	80
4.7.2	Out-of-sample Comparison	80

4.8	Summary	83
5	RMT for Foreign Exchange Portfolios	87
5.1	Introduction	87
5.2	Data	89
5.3	Accuracy of the Limiting Approximations for 39 Assets	89
5.4	Measuring Noise	91
5.5	In-sample Testing	91
5.5.1	Methodology	91
5.5.2	Further Reduction of S&P 500 Asset Numbers	93
5.5.3	In-sample Analysis of the Fx Portfolio with 39 Assets	95
5.5.4	Further Reduction of Fx Asset Numbers	97
5.6	Out-of-sample Testing	102
5.6.1	Methodology	102
5.6.2	Out-of-sample Analysis	102
5.6.2.1	Overall Risk	102
5.6.2.2	Annual Risk	103
5.6.2.3	Daily Risk	104
5.6.3	Parameter Values	105
5.6.4	Range of Realised Risks	107
5.6.5	Reduction of Assets	111
5.7	Summary	112
6	Conclusions and Future Work	115
6.1	Goals of the Thesis	115
6.2	Conclusions	116
6.2.1	Overview	116
6.2.2	Stability Filtering	117
6.2.3	Foreign Exchange Filtering	118

6.2.4	Choice of Weighting System and Parametrisation	120
6.2.5	Covariance Filtering	121
6.2.6	Summary	121
6.3	Future Work	122
References		124
A Maximum Eigenvalue of an Exponentially Weighted Random Matrix		1
A.1	Introduction	1
A.2	Maximum Eigenvalue of an Exponentially Weighted Random Matrix	2
A.3	Summary	5
B List of Publications		1
C Matlab[®] source code disc		1

Abstract

This thesis describes research on filtering methods using Random Matrix Theory (RMT) Models in financial markets. In particular, a novel, stability-based RMT filter is proposed and its potential, for reducing stock portfolio risk, is compared to two well-known alternatives. In terms of performance, the stability-based filter achieved 17.3% overall improvement in risk reduction for equally weighted forecasts, and 49.2% for exponentially weighted. Of the filters investigated, not only did it prove to be the most effective and consistent, for overall risk reduction, but was also shown to reduce the frequency of large risk increases, (which, despite their importance, have attracted little attention in the literature to date). The full frequency distribution of filter effects is studied and a comprehensive test methodology established. Improvements, on previous approaches, include integrated use of bootstrap analysis and out-of-sample testing. RMT filtering was also applied to the foreign exchange market, which contains far fewer assets than a typical stock portfolio. Filters were shown to reduce inherent currency trading risks, despite the small number of assets involved. Once again, our novel filter resulted in the lowest risk for exponentially weighted forecasts, and was most consistent in reducing overall levels, exhibiting also the fewest large risk increases. Finally, and more generally, RMT filter testing and analysis can be used to demonstrate the value of rapid response models, i.e. those reacting quickly to market events. Despite the fact that these utilise very recent data, much information is typically masked by noise. Filtering is shown to be successful in exposing such key underlying features.

Acknowledgements

I would like to thank the Irish Research Council for Science, Engineering and Technology for supporting this research under the Embark Scholarship Scheme. I would also like to thank all my friends and family, of course.

Finally, I would like to offer special thanks to both of my supervisors, Prof. Heather J. Ruskin and Dr. Martin Crane, for their guidance and encouragement throughout this work.

Thank you all.

Justin Daly,

17 June, 2009.

List of Tables

3.1	List of KR methods tested	35
3.2	List of in-sample scenarios used to test the effect of filtering on risk, and where the main results are reported, for different portfolio sizes (N). Three different categories have been tested: (1) filtering the S&P correlation, (2) filtering the S&P covariance, and (3) filtering the Fx correlation.	43
3.3	List of out-of-sample scenarios used to test the effect of filtering on risk, and where the main results are reported, for different portfolio sizes (N). Three different categories have been tested: (1) filtering the S&P correlation, (2) filtering the S&P covariance, and (3) filtering the Fx correlation. . .	44
4.1	Mean out-of-sample realised risk as a percentage of that for unfiltered equally weighted covariance, where correlation was filtered. Filtering was seen to reduce mean realised risk in all cases, and to as low as 84% of our benchmark risk, when applying KR filtering to exponentially weighted forecasts.	60
4.2	Mean out-of-sample realised risk per year for 6 years as a percentage of the equally weighted unfiltered result each year, filtering correlation. In some cases, RMT filtering was seen to increase mean realised risk over the course of a year. The KR filter was the only one to reduce realised risk in all cases each year.	60

4.3	Daily frequency of percentage effect, on realised risk, of applying RMT filters to the correlation matrix. “Low” and “High” specify a range for the realised risk, expressed as a percentage of the relevant unfiltered realised risk. Filtering was seen to reduce risk 74.3% of the time, while stability-based filtering reduced it the most often, 80.7% of the time. It can also be seen that, despite reducing mean realised risk overall, all filters had the potential to markedly increase realised risk on any individual day.	61
4.4	Summary of filter performance in different unfiltered risk environments, for the S&P 500 portfolio, for equally weighted covariance forecasts. “All data” refers to the full range of unfiltered realised risk. The remaining columns break this range into quintiles. The results shown are the average realised risk in each quintile, as a percentage of the corresponding unfiltered result. Thus it can be seen that the filters were most effective at reducing risk during periods where that risk was high. This resulted in the range of risk being improved upon the unfiltered case for all filters.	67
4.5	Summary of filter performance in different unfiltered risk environments, for the S&P 500 portfolio, for exponentially weighted covariance forecasts. This table corresponds to Table 4.4. It can again be seen that the filters were most effective at reducing risk during periods where that risk was high. In this case, this resulted in the range of risk being equivalent to (PG+) or improved upon (LCPB, KR) the unfiltered range.	67
4.6	Mean out-of-sample realised risk as a percentage of that for unfiltered equally weighted covariance, with direct filtering of the covariance matrix. Filtering was seen to reduce mean realised risk in all cases, and to as low as 86.7% of the benchmark, for the KR filter applied to exponentially weighted covariance.	74

4.7	Mean out-of-sample realised risk per year for 6 years as a percentage of equally weighted unfiltered result each year, filtering covariance. In some cases, RMT filtering was seen to increase mean realised risk over the course of a year.	75
4.8	Daily frequency of percentage effect, on realised risk, of applying RMT filters directly to the covariance matrix. “Low” and “High” specify a range for the realised risk, expressed as a percentage of the relevant unfiltered realised risk. Filtering was seen to reduce risk 70.3% of the time, while stability-based filtering reduced it the most often, 80.4% of the time. It can again be seen that, despite reducing mean realised risk overall, all filters had the potential to markedly increase realised risk on any individual day. .	75
4.9	Summary of filter performance in different unfiltered risk environments, for the S&P 500 portfolio, for equally weighted covariance forecasts, using direct covariance filtering. “All data” refers to the full range of unfiltered realised risk. The remaining columns break this range into quintiles. The results shown are the average realised risk in each quintile, as a percentage of the corresponding unfiltered result. Thus it can be seen that the filters were most effective at reducing risk during periods where that risk was high. .	79
4.10	Summary of filter performance in different unfiltered risk environments, for the S&P 500 portfolio, for exponentially weighted covariance forecasts, using direct covariance filtering. This table corresponds to Table 4.9. We again see that the filters were most effective at reducing risk during periods where that risk was high.	79
4.11	Percentage of the time each method had the lowest mean out-of-sample realised risk on an annual, monthly and daily basis. The most consistent method by far was the stability-based filter, applied to the correlation matrix of the exponentially weighted forecasts.	82

4.12	Mean out-of-sample realised risk as a percentage of that for unfiltered equally weighted covariance, for all methods, and sorted best to worst. The three best methods involved KR filtering, while four of the best five used exponential weights.	83
5.1	List of currencies and commodities used. The base currency was the U.S. Dollar.	88
5.2	Mean out-of-sample realised risk as a percentage of that for unfiltered equally weighted covariance, for the full Fx portfolio, with 39 assets. Filtering was seen to reduce mean realised risk in all cases.	103
5.3	Mean out-of-sample realised risk per year, for 7 years, as a percentage of the equally weighted unfiltered result each year, for the full Fx portfolio, with 39 assets. RMT filtering was seen to increase mean realised risk over the course of a year in a few cases, and for all filters, and to reduce it in the majority of cases.	104
5.4	Daily frequency, of percentage effect on realised risk, of applying RMT filters to the Fx portfolio. “Low” and “High” specify a range for the realised risk, expressed as a percentage of the relevant unfiltered realised risk. Similar to the S&P 500 case, RMT filtering was seen to reduce realised risk in the majority of cases, while also having the potential to increase realised risk, more than doubling it on some individual days.	105
5.5	Summary of filter performance in different unfiltered risk environments, for equally weighted covariance forecasts. “All data” refers to the full range of unfiltered realised risk. The remaining columns break this range into quintiles. The results shown are the average realised risk in each quintile, as a percentage of the corresponding unfiltered result. Thus it can be seen that the filters were most effective at reducing risk during periods where that risk was high.	108

5.6	Summary of filter performance in different unfiltered risk environments, for exponentially weighted covariance forecasts. This table corresponds to Table 5.5. It can again be seen that the filters were most effective at reducing risk during periods where that risk was high.	108
5.7	Mean out-of-sample realised risk as a percentage of that for unfiltered equally weighted covariance, for the Fx portfolio with 15 major assets. Risk was increased after filtering in all but one case.	111
5.8	Mean out-of-sample realised risk per year, for 7 years, as a percentage of the equally weighted unfiltered result each year, for the Fx portfolio with 15 major assets. Compared to the 39 assets case, we saw many more years with increases after filtering, and larger increases were observed.	111
5.9	Daily frequency, of percentage effect on realised risk, of applying RMT filters to the Fx portfolio with 15 major assets. “Low” and “High” specify a range for the realised risk, expressed as a percentage of the relevant unfiltered realised risk. The range of frequencies was seen to be narrower than that of the full Fx Portfolio, while a higher frequency of risk increases was observed.	112
6.1	Mean realised risk as a percentage of unfiltered, for all filters and both portfolios. The improvements show remarkable consistency between the two asset classes.	119

List of Figures

4.1	Percentage of eigenvalues that were larger than the maximum eigenvalue predicted by RMT, for equally weighted correlation and covariance with 100 assets (above), and 432 assets (below).	50
4.2	Percentage of eigenvalues that were larger than the maximum eigenvalue predicted by RMT, for exponentially weighted correlation and covariance with 100 assets (above) and 432 assets (below).	51
4.3	In-sample mean stability for the equally weighted covariance forecast with 100 assets, filtering correlation.	52
4.4	In-sample mean stability for the equally weighted covariance forecast with 100 assets, filtering covariance.	53
4.5	Mean bootstrapped (in-sample) realised risk, for selected filters, applied to equally weighted volatility forecasts of S&P 500 stocks, and for unfiltered volatility (“ORIG”), for 100 assets. The x-axis shows the Q -value, which is the number of past moves used, divided by the number of assets in the portfolio. Note that filtering is most effective for low values of Q , and produces the best overall risk.	54

4.6	Mean bootstrapped (in-sample) realised risk, for selected filters, applied to equally weighted volatility forecasts of S&P 500 stocks, and for unfiltered volatility (“ORIG”), for 250 assets (above) and 432 assets (below). The x-axes show the Q -value, which is the number of past moves used, divided by the number of assets in the portfolio. Filters were more effective, overall, for the larger number of assets, and the best risk occurred, with stability filtering, at the lowest possible value of Q in both cases. Note also the consistency between the stability filters shown.	55
4.7	Mean bootstrapped (in-sample) realised risk, for selected filters, applied to exponentially weighted volatility forecasts of S&P 500 stocks, and for unfiltered volatility (“ORIG”), for 100 assets (above), and with detailed view (below). RMT filtering is seen to improve risk across a wide range of decay factors, with a stability filter (KR2) resulting in the best overall risk. .	56
4.8	Mean bootstrapped (in-sample) realised risk, for selected filters, applied to exponentially weighted volatility forecasts of S&P 500 stocks, and for unfiltered volatility (“ORIG”), for 250 assets (above), and with detailed view (below). The lowest risk was again produced by using a stability filter.	57
4.9	Mean bootstrapped (in-sample) realised risk, for selected filters, applied to exponentially weighted volatility forecasts of S&P 500 stocks, and for unfiltered volatility (“ORIG”), for 432 assets (above), and with detailed view (below). Stability filters can again be seen to outperform other models, while two local minima for the decay factor are also seen. Note that the Riskmetrics decay factor of 0.97 performed poorly without filtering (for all three values of N).	58
4.10	Optimal forward validation number of past moves for equally weighted covariance, filtering the correlation matrix. Note that low values of T were usually chosen for all three filters.	63

4.11	Optimal forward validation decay factor values for filtering the correlation matrix. The range of decay factors seen here was higher in all cases than the Riskmetrics (1996) value of 0.97.	63
4.12	First thousand normalised exponential weights corresponding to decay factors of 0.97 and 0.999. Note that the style of models generated are markedly different, with the weights corresponding to 0.999 behaving almost like an equally weighted model.	64
4.13	Filtered vs. unfiltered realised risks, for the S&P 500 portfolio, for equally (above) and exponentially (below) weighted forecasts. The unfiltered risk is also included, as a 45 degree line, for comparison. It can be seen that filtering was most effective when the unfiltered series was in the upper end of its range, while the large risk increases occurred in the middle and lower ends of this range. Overall the upper limit of the <i>range</i> of the filtered risks was either equivalent to, or improved upon, that for the unfiltered risk. . . .	66
4.14	Mean bootstrapped (in-sample) realised risk, for selected filters, filtering covariance, with exponentially weighted volatility forecasts of S&P 500 stocks, and for unfiltered volatility (“ORIG”), for 100 assets (above), and with detailed view (below). RMT filtering is again seen to be effective across a wide range of decay factors, with stability filtering being preferred. . . .	69
4.15	Mean bootstrapped (in-sample) realised risk, for selected filters, filtering covariance, with exponentially weighted volatility forecasts of S&P 500 stocks, and for unfiltered volatility (“ORIG”), for 250 assets (above), and with detailed view (below). Stability filtering is again preferred in-sample, with good consistency between the stability filters shown.	70

4.16	Mean bootstrapped (in-sample) realised risk, for selected filters, filtering covariance, with exponentially weighted volatility forecasts of S&P 500 stocks, and for unfiltered volatility (“ORIG”), for 432 assets (above), and with detailed view (below). These figures have much in common with those for filtering correlation, with stability filtering again found to result in the best risks. Here though we see one clear minimum, and subsequent optimal decay factor.	71
4.17	Mean bootstrapped (in-sample) realised risk, for selected filters, filtering covariance, with equally weighted volatility forecasts of S&P 500 stocks, and for unfiltered volatility (“ORIG”), for 100 assets (above) and 250 assets (below). The x-axes show the Q -value, which is the number of past moves used, divided by the number of assets in the portfolio.	72
4.18	Mean bootstrapped (in-sample) realised risk, for selected filters, filtering covariance, with equally weighted volatility forecasts of S&P 500 stocks, and for unfiltered volatility (“ORIG”), for 432 assets. In this case, an overall preference for the lowest value of Q , with filtering, can again be seen, similar to the correlation case.	73
4.19	Optimal number of past moves for equally weighted covariance, filtering the covariance matrix directly. Here we see that the best parameter choices varied across the range, over time.	76
4.20	Optimal forward validation decay factor values for filtering the covariance matrix directly. The range of decay factors seen here was higher in all cases than the Riskmetrics (1996) value of 0.97.	77

4.21	Filtered vs. unfiltered realised risks, for the S&P 500 portfolio, for equally (above) and exponentially (below) weighted forecasts, using direct covariance filtering. The unfiltered risk is also included, as a 45 degree line, for comparison. It can be seen that filtering was most effective when the unfiltered series was in the upper end of its range, while the large risk increases occurred in the middle and lower ends of this range. For equal weights, the upper limit of the <i>range</i> of the filtered risks was less than that for the unfiltered risk, while, for exponential, some increase in the upper limit was seen.	78
4.22	In-sample mean realised risk for selected best forecasting methods, for 100 assets (above) and 432 assets (below), showing both equal and exponential weighting schemes. Stability filters are preferred in all cases.	81
5.1	Eigenvalue distribution of 1000 sampled random matrices vs RMT approximation, for 39 assets, with equal weights and $T = 40$ historical data points (above), and with exponential weights and a decay factor of $\alpha = 0.97$ (below).	90
5.2	Percentage of eigenvalues that were larger than the maximum eigenvalue predicted by RMT, for equally (above) and exponentially (below) weighted covariance, with 39 assets. As the amount of past data is increased, (corresponding to an increase of decay factor in the exponential case), the amount of non-noisy eigenvalues is increased. However, using long price histories may not lead to the best medium term forecasts.	92

5.3	Mean bootstrapped (in-sample) realised risk, for selected filters, applied to equally (above) and exponentially (below) weighted volatility forecasts of S&P 500 stocks, and for unfiltered volatility (“ORIG”), for 39 assets. For equal weights, the x-axis shows the Q -value, which is the number of past moves used, divided by the number of assets in the portfolio. Note that the trend of the 432, 250 and 100 asset cases (Figures 4.5 to 4.9) was continued here, namely that the unfiltered forecasts became more competitive as the number of assets was reduced.	94
5.4	Mean bootstrapped (in-sample) realised risk, for selected filters, applied to equally (above) and exponentially (below) weighted volatility forecasts, and for unfiltered volatility (“ORIG”), for the Fx portfolio, with 39 assets. Filtering is seen to have been effective in both cases. For equal weights, a preference for low Q -values, with filtering, can be seen, and the closest agreement is with the S&P 500 case with $N = 432$, (Figure 4.6). In the exponential case, with filtering, the optimal decay factor (0.97) coincided with Riskmetrics (1996).	96
5.5	Mean bootstrapped (in-sample) realised risk, for selected filters, applied to equally weighted volatility forecasts of the Fx portfolio, and for unfiltered volatility (“ORIG”), for 30 (above) and 20 (below) assets. Filtering is still effective, reducing overall risk in both cases.	98
5.6	Mean bootstrapped (in-sample) realised risk, for selected filters, applied to equally weighted volatility forecasts of the Fx portfolio, and for unfiltered volatility (“ORIG”), for 15 (above) and 10 (below) assets. Note that as N was reduced, the unfiltered forecasts became more competitive.	99
5.7	Mean bootstrapped (in-sample) realised risk, for selected filters, applied to exponentially weighted volatility forecasts of the Fx portfolio, and for unfiltered volatility (“ORIG”), for 30 (above) and 20 (below) assets. Stability filtering is still effective in both cases.	100

5.8	Mean bootstrapped (in-sample) realised risk, for selected filters, applied to exponentially weighted volatility forecasts of the Fx portfolio, and for unfiltered volatility (“ORIG”), for 15 (above) and 10 (below) assets. As for equal weights, the unfiltered forecast was shown to be more competitive as N was reduced.	101
5.9	Optimal parameters for equally (above) and exponentially (below) weighted forecasts, as selected by forward validation, for the Fx portfolio. In the unfiltered cases (not shown), the number of past moves ranged between 200 and 240, and the decay factors were usually equal to 0.98, and sometimes in the range 0.99 - 0.995. A tendency toward reactive models can be seen in these graphs.	106
5.10	Filtered vs. unfiltered realised risks, for equally (above) and exponentially (below) weighted forecasts. The unfiltered risk is also included, as a 45 degree line, for comparison. It can be seen that filtering was most effective when the unfiltered series was in the upper end of its range, while the largest percentage risk increases occurred in the middle and lower ends of this range. Overall, the upper limit of the range of the filtered risks was less than that for the unfiltered risk in both cases.	109
5.11	Forecast and realised risks, for unfiltered (above) and KR filtered (below) equally weighted covariance forecasts, over the life of the forward validation. We see a shortfall in forecasting in the middle and, particularly, upper sections of the range. This “flat” forecast profile, and resulting shortfall, was common to all three filters, (and the unfiltered), and to both types of weighting scheme. This highlights the difficulty in forecasting financial markets.	110
A.1	$x - \ln(x)$	5

Chapter 1

Introduction

1.1 Motivation

Over the years, research into financial markets has borrowed from many other fields, and investment professionals now use a wide range of tools in an effort to maximise their returns and minimise their risks. Random Matrix Theory (RMT) is one such technique, originating in statistical physics, which has recently been applied to studying financial correlations. RMT has already been used successfully to remove noise in these correlation matrices, leading to improved investment choices. Much of the focus for RMT applications to date has been the analysis of the stock market, with its large asset class, but this represents only one aspect of practical trading, which often deals with the management of smaller asset groups. In what follows, we look at ways of augmenting RMT filters and examine their potential to improve foreign exchange portfolios, one example of a class with limited assets.

From the perspective of RMT, we review research which outlines how the eigenvalues, of the correlation matrix of stock market returns¹, are consistent with those calculated using random returns, with the exception of a few large eigenvalues. Moreover, filtering techniques based on eliminating these “noisy” eigenvalues have been beneficial in both re-

¹in this case we employ a daily return, which is the log of todays price divided by yesterdays, as defined in Section 3.2.3

ducing the realised risk² of optimised portfolios, and improving the forecast of this realised risk. Similar results have also been measured in other markets, such as volatility markets³, and even fund of hedge fund management⁴.

In this thesis, we have assessed three such RMT filters, from the perspective of a portfolio manager with a medium-term investment horizon, and an active (managed) investment mandate. We tested how these filters can be used to reduce the realised risks of traded portfolios, and what sort of behaviour can be expected from them on a daily basis in the trading room. With this we aimed to provide a basis from which these filters can be integrated into different trading applications.

1.2 Objectives

Our main objectives in this research were as follows.

1. Following the work of Sharifi et al. (2004), we proposed to test a novel RMT filter, which we have developed, based on improving the stability of the filtered matrices. The original filter was constructed in the context of high frequency data, and was found to improve stability. We have modified and enhanced it, and applied it to our daily data. We then compared this filter with two well known RMT filters from the point of view of stability, risk reduction, and consistency, and assessed the behaviour of all three filters in the stock market, and then foreign exchange.
2. We examined the use of RMT filters in the foreign exchange market. The stock market, a typical market for testing these filters, is characterised in this instance by the low availability of historical data *in comparison with the number of tradeable assets*. It is this ratio that affects the amount of noise in the risk forecasts and, in the case of the stock market, the ratio is low. This lends itself to a high expected

²the realised risk is the standard deviation of the portfolio return over the investment period, as defined in Section 3.3.1

³A volatility trade is one where the asset traded is the standard deviation of some other asset.

⁴A fund of funds manager is one who invests in a portfolio of *funds*, using many of the same portfolio management techniques used to invest in the underlying assets.

level of noise. In contrast, the foreign exchange market has relatively few assets. However, an auxiliary objective was to test the view that, just because historical data are available, these are not necessarily useful. Thus the availability of *relevant* past data is in question, and RMT filters may prove to be valuable in practice.

3. At the same time, we wished to build and improve upon the existing test methodology available in the literature on RMT filtering. Firstly, we tested the filters using forward validation, a simulation of a live environment, which can have no knowledge of future events when performing either model parametrisation, or evaluation. Also, while work has been done on the aggregate results of RMT filtering, we were keen to assess the details. We aimed to evaluate both the annual and day-to-day performance of these filters, to gain a deeper understanding of their behaviours than was previously available. This is a necessary step for any potential live trading application, since the trading room should demand to be appraised of typical low level expectations in practice.
4. We tailored these tests for the type of problems that we were interested in. Following on from prior industry experience, we were primarily concerned with medium term investments, and have used 20 days as our main forecast period to coincide with the work of Pafka et al. (2004). We were focused on the role of an active manager, in which the expectation is that trading will take place on a daily basis. This is in contrast to a passive investment, which is reviewed quarterly, for example.
5. We wished to evaluate a close approximation to a live trading application, while making a limited number of underlying assumptions about the trading strategy. In this way we planned for our research to be of benefit to a wide range of applications. To achieve this we used the solutions to the minimum risk problem as our test portfolios.
6. We also aimed to assess the choice of weighting method in the forecasts, along with the best parameters for each weighting, with and without filtering. We have considered two alternative weighting systems in the thesis, equal and exponential. Both of

these can be filtered using RMT.

7. Finally, we wanted to study the effect of filtering the covariance matrix directly. The standard implementation is to filter correlation, which is closer to the RMT assumptions. However, the covariance matrix should contain more information and so we considered direct filtering of it here.

1.3 Thesis Outline

The thesis has been organised in the following way. In Chapter 2 we review the related literature, showing how RMT has been applied in the past to improving the optimisation of financial portfolios, and we discuss the different filters that have been tested. In Chapter 3 we define more strictly the concepts used in the thesis, and we detail the test methodologies that have been applied. Here, we also define our novel, stability based RMT filter.

The results have been divided into two chapters, based on the market the filters were applied to. In Chapter 4 we study the application of RMT filtering to an S&P 500 portfolio. We first assess the amount of noise in the system, as measured by RMT. Next, we examine the effect of RMT filtering on stability. We then compare the effect of the filters on in-sample risk, (using bootstrapping), for three different portfolios sizes, and compare our work with existing analysis. We then study the filters out-of-sample, using forward validation. Here, we examine them also on an annual, monthly, and daily basis. We also compare the range of the realised risk with and without filtering, and study the behaviour of the filters as a function of the unfiltered risk. This was motivated by the daily analysis, which revealed that filtering did not always decrease risk. Both in and out-of-sample, we consider the optimal model parameters, and compare these to previous results. We also examine the application of the filters to the covariance matrix, which is less directly related to the assumptions of RMT, but potentially holds more information.

In Chapter 5 we apply the filters to the foreign exchange market. Here we consider a portfolio with only 39 assets. We provide a comparative analysis to the stock portfolio case,

again studying the amount of noise in the system and the effect of filtering on risk, using bootstrap analysis and forward validation. Risk is once more evaluated annually, and daily, and the range of filtered risk is also considered. We also assess the optimal choice of model parameters, for both weighting systems, and how this choice is affected by filtering. Finally, we consider the limitations of the filters, by further reducing the number of tradeable assets. We compare this to the effect of further asset reduction in the stock market.

In Chapter 6 we summarise our work. We discuss our conclusions, from the point of view of a practitioner wishing to implement the technology. We also discuss the future research that would be required to apply the filters in a hedge fund context, (involving integration with other elements of a mathematical trading system), as well as potential further improvements to the filters themselves.

In Appendix A, we review a numerical technique we have developed, for improving the runtime needed to filter exponentially weighted risk forecasts using RMT. Appendix B contains a list of publications related to this research. Finally, a disc of the Matlab[®] source code can be found in Appendix C.

Chapter 2

Literature Review

2.1 Introduction

In recent years, the availability of large amounts of financial market data has provided an opportunity for researchers from a wide range of fields to contribute to the study of financial markets. Econophysics is one such field, which has developed a large body of work. On reviewing this work, much attention has been paid to Bouchaud and Potters (2000, 2003), where developments in areas such as the distributions of prices, the correlations within and between assets, extreme events, and portfolio management, are discussed. A similar overview is presented in Mantegna and Stanley (2000).

In Bouchaud and Potters (2000, 2003) the focus was on the interpretation of financial data, using concepts from statistical physics in particular. One aspect studied was distributional behaviour. Despite many theories of finance assuming normality of log returns, it has been known by market practitioners for some time that these models have their limitations, (Hull, 2006). For example, options are priced using volatility “smiles” which effectively compensate for non-normality. In Bouchaud and Potters (2003), the authors concluded that price returns more appropriately could be fitted by truncated Lévy or Student distributions. Moreover, the tails could be fitted by a Pareto (power-law) tail. Risk was also investigated, and complex fluctuations in volatility were noted, along with temporal

correlations in the volatility time series. The importance of the correlations between assets, central to this thesis, is further discussed below. Here, Bouchaud and Potters (2000, 2003) showed good agreement between the eigenstates of market correlations, and a random system. Other areas investigated included extreme events, portfolio theory, and option pricing. What is most notable about this body of work is its efforts to remain close to practical questions encountered in trading rooms by financial engineers.

This applied perspective can also be found in other areas of econophysics. For example, Dacorogna et al. (2001) have studied high frequency effects, such as stylized facts, seasonality, volatility and correlation. One result, particularly relevant to this thesis, is that the behaviours of high frequency and daily volatility were quite different, and simply scaling between them was not recommended. We have concentrated solely on daily data here, and have considered a 20 and 50 day forecasting period. Thus caution should be exercised in making any predictions about higher frequency effects based on this work.

Gençay et al. (2002) described how wavelets can be used to examine financial data, in particular seasonality, to identify structural breaks, to study behaviours at different timescales, and to remove high frequency effects from data prior to forecasting. This latter point is an interesting area for future research, which may complement RMT filtering, perhaps by revealing additional non-randomness. Meanwhile, many other types of filter have also been applied in a financial context. A wide variety of these can be found, for example, in Brigo and Hanzon (1998), Gençay et al. (2002), Tsay (2005) and Tumminello et al. (2005).

Much of the above analysis has aimed to understand the net behaviour of traded assets, but recent research has also examined the trader viewpoint through agent based modelling (e.g. Bak et al. (1997); Sornette and Johansen (1998); Zhu et al. (2009) and references therein). Such models aim to describe the behaviours of individual agents in the market. Following from this, it is possible to study the net result of these behaviours and compare this to actual data, therefore improving our understanding of the drivers of price movement.

In the remainder of this chapter we review the key work which has underpinned this thesis. We begin by reviewing traditional Markowitz portfolio theory, followed by the fore-

casting of financial covariances, one of the key unknown parameters of that theory. We then review the role of RMT in improving Markowitz portfolio optimisation. Central to this thesis is the work of three groups: Laloux et al. (1999; 2000); Plerou et al. (1999; 2000a; 2000b; 2001; 2002); and Pafka and Kondor (2002a,b), Pafka et al. (2004). We review some of this in detail, together with related work in the field of RMT in finance. In this context, Conlon et al. (2007), discusses the application of RMT to a market with a small number of assets, and this has informed our own study on foreign exchange filtering.

We also discuss the work of Sharifi et al. (2004), which introduced the concept of stability-based filtering into the RMT literature, and from which we have developed an extension of that filter, tested here. Finally, we look at some ways in which eigenanalysis has been used to provide further information about market activity.

2.2 Portfolio Theory and Noise

Markowitz (1959) portfolio theory is an intrinsic part of modern financial analysis. The aim of portfolio theory is to determine the optimal portfolio weights which either maximise return for a fixed level of risk, or minimise risk for a fixed return. Constraints on the optimisation can be imposed, which reflect the problem at hand, although care should be taken, since certain formulations for constraints may cause multiple and unstable optimal solutions, (Galluccio et al., 1998; Bongini et al., 2002).

In all its formats, portfolio theory relies on the covariance matrix of returns. However, this can be difficult to estimate. For example, for a time series of length T , a portfolio of N assets requires $(N^2 + N)/2$ covariances to be estimated from NT returns. This results in estimation noise, since the availability of historical information is limited. Moreover, it is commonly accepted that financial covariances are not fixed over time (e.g. Morgan and Reuters (1996); Dacorogna et al. (2001); Bouchaud and Potters (2003)); thus older historical data, even if available, can lead to cumulative noise effects.

In this thesis, we consider two methods of forecasting the covariance matrix, namely

equally and exponentially weighted. We now review these methods for unfiltered forecasting, and later discuss the filtering methods for each one.

2.3 Volatility Forecasting and Riskmetrics

Market covariances vary over time, having periods of sustained low or high values (Morgan and Reuters (1996); Dacorogna et al. (2001); Bouchaud and Potters (2003)). Of the two covariance forecasting methods considered in this thesis, equally and exponentially weighted, the former are extensively discussed in the literature on RMT, (Section 2.4). This method consists of a simple moving average of past data. Exponentially weighted forecasts meanwhile, give more weighting to recent events, and use an exponential decay factor to discount the effect of older data. The definitions of these methods are given in Chapter 3.

Exponentially weighted covariance forecasting was popularised by, for example, J.P. Morgans Riskmetrics (1996) work. It is now considered a benchmark in risk management, (Pafka et al., 2004). This method is held to have certain advantages over the equally weighted model, that are amplified when there is a shock in the market, (Riskmetrics, 1996). Firstly, exponentially weighted models react faster to new market information, since more recent data is more heavily weighted than older data. At the same time, the effect of an event is slowly removed as time moves forward, by use of the exponential decay factor.

In contrast, equally weighted forecasts do not give preference to recent events by definition. All events in the system are given the same weighting, no matter when they occurred, as long as they are within the time period considered. Moreover, when an event falls outside of this time period covariance can change suddenly which, in conceptual terms, is very unsatisfactory.

In the Riskmetrics (1996) work, optimal decay factors were calculated for daily and monthly forecasting. The recommended values were 0.94 (daily) and 0.97 (monthly). Other sources have suggested the use of higher values, for example Pafka et al. (2004) recommended, using RMT analysis, a value of 0.996 for one month forecasting of a large stock

portfolio, while Litterman and Winkelmann (1998) analysed a bond portfolio, using maximum likelihood estimation, and reported values of 0.98 (daily) and 0.99 (monthly). In this thesis, we have examined the optimal choice of decay factor for our particular problem.

Riskmetrics (1996) also discussed the heteroskedasticity of financial markets, using the example of exchange rates and the S&P 500, where periods of low volatility and periods of high volatility were clearly seen in both cases. Covariance was seen to display similar behaviour. This work is a good example of why the use of long time series of historical data for medium term forecasting is counter-intuitive.

When forecasting the covariance matrix, Riskmetrics (1996) considered a number of alternative methods, including the use of *implied volatility*¹, which has a number of drawbacks. For example, it depends on the option pricing model used, while most option models assume standard deviation is constant. Also, implied volatility is associated with fixed intervals of time (1 week, 1 month, 3 months etc.) and so intervening times require some type of interpolation. In general, the availability of data may be problematic, and implied correlations in particular are hard to source. Covariance matrices calculated from implied volatility are also not guaranteed to be positive definite.

Instead, exponentially weighted covariance forecasts were suggested, and the question of optimal decay factors was addressed. While different decay factors are, in theory, possible for different assets, the problem quickly becomes very complex as the number of assets increases. This is due to maintaining the properties of the covariance matrix, such as non-negative variances, symmetry, and correlations in the appropriate range. Therefore, the use of a single decay factor across all assets was recommended for use with Riskmetrics (1996).

However, different decay factors were used when considering different forecasting periods, since simply scaling up shorter-term estimates is inappropriate. For example a short term rise in the level of volatility will result in short term forecasts being higher. But this cannot be scaled out over a long time period without overestimating risk, since this short term rise will not be sustained. Finally, the optimal decay factors themselves were calcu-

¹Implied volatility is volatility calculated from traded option prices, by inverting an option pricing formula.

lated using a root mean-squared error approach, where the individual errors were taken to be the differences between the forecast and realised risk.

In this thesis, Riskmetrics style forecasts, and parameters, will be tested alongside forecasts made using equal weights, and also filtered forecasts using both weighting systems. We now discuss how RMT has been used to filter both types of forecast.

2.4 Applications of Some Aspects of Random Matrix Theory to Portfolio Selection

Random matrix theory (RMT) was first developed by authors such as Dyson and Mehta (1962; 1963; 1963; 1991), to explain the energy levels of complex nuclei (Plerou et al., 2002). It was recently applied (by several authors including Plerou et al. (1999; 2000a; 2000b; 2001; 2002) and Laloux et al. (1999; 2000)) to noise filtering in financial time series, particularly in large dimensional systems such as stock markets. We now review the work of these two groups, who applied RMT to equally weighted covariance.

2.4.1 Work of Laloux et al. (1999; 2000)

Laloux et al. (1999; 2000) were one of the first groups to research the application of RMT to finance, addressing the similarities between the eigensystems of financial correlation matrices, and those of random matrices, and, importantly, where the differences lay. These authors, primarily, examined an S&P 500 dataset, involving daily data over the period 1991 to 1996. A total of 406 stocks covered the entire time interval. On comparing the statistics, of the eigenvalues and eigenvectors of the market correlations, to those of a corresponding random matrix, Laloux et al. (1999, 2000) reported the following key observations.

- The largest eigenvalue was found to be 25 times larger than the maximum eigenvalue predicted by RMT.
- On comparing the remaining eigenvalues with RMT, good agreement was found between the distributions, after fitting the total risk in the system to an appropriate value,

in effect compensating for any further “non-noisy” eigenvalues.

- In this way they determined that 94% of the eigenvalues were within the noise band predicted by RMT, while the highest 6% of eigenvalues were above the maximum random eigenvalue. These largest 6% of eigenvalues were found to contain 26% of the total system volatility.
- An examination of the eigenvectors themselves revealed that those, corresponding to the eigenvalues which were within the noise band, had components which were consistent with randomness. At the same time, the eigenvector for the largest eigenvalue was shown to have clearly non-random elements. They concluded that this eigenvector represented the “market”, in the sense that it assigned a roughly equal weighting to each stock.
- The “non-noisy” eigenvectors were found to be more stable in time.
- This analysis was also applied to other markets, including volatility markets, with similar results.

The authors thus concluded that the use of such correlation matrices, which corresponded so closely to pure noise, for portfolio optimisation, was unwise. Instead they suggested a filtering algorithm based on RMT (Laloux et al., 2000) to isolate the non-noisy information.

It should be noted that, in using such RMT filtering, *empirical* analysis is performed. The observed statistics, of the bulk of the eigenvalues are, repeatedly across different markets, found to correspond to the statistics of random matrices, (further discussed in this chapter). Following Laloux et al. (2000) this can be explained by assuming that the components of the correlation matrix which are orthogonal to the eigenvectors of the “large” eigenvalues is pure noise. This approach to market analysis differs from other approaches that can be taken in financial econometrics and financial mathematics where the model can be based on, for example, the capital asset pricing model or asset pricing theory, and

the noise is usually taken to be stochastically independent of the “signal”. Such models lend themselves more easily to techniques such as Bayesian filtering (e.g. Brigo & Hanzon (1998)).

In Laloux et al. (2000), filtering correlation matrices using such RMT analysis was found to improve portfolio optimisation, as follows. First a date in the centre of the data was chosen, and the data was split into “past” and “future” periods. We note that this future period involved a fixed long term investment, over many years, unlike our medium term (monthly) investments. The past data was used to calculate a correlation matrix, which was then filtered. The filter described by Laloux et al. (2000) works by assuming that the market eigenvalues below the maximum random eigenvalue (defined Section 3.2.3) are noise. These noisy market eigenvalues are then filtered out, by replacing them with constant values, while maintaining the trace of the system (i.e. the sum of the eigenvalues) so that the system does not become distorted.

Using both the filtered and unfiltered matrices, they constructed two efficient frontiers resulting from Markowitz optimisation. In doing this, the correlation matrix was the focus, and future returns were assumed known when calculating the variance of the individual assets. Once the efficient frontier portfolios were known, the authors were then able to calculate the realised risk of these portfolios over the “future” period. They established the following key results, which are of particular relevance to this thesis.

- Portfolios based on the unfiltered correlation greatly underestimated the actual realised risk in the future period, at times by a factor of three.
- While the filtered results also underestimated realised risk, the forecast error was greatly reduced, to a factor of order 1.5.
- The realised risk of the filtered portfolios was below that for the unfiltered at all points on the efficient frontier. In fact, the difference in filtered and unfiltered risk was seen to be fairly consistent along the frontier.

We now examine the work of a second group, Plerou et al. (1999; 2000a; 2000b; 2001; 2002) which was undertaken over a similar period.

2.4.2 Work of Plerou et al. (1999; 2000a; 2000b; 2001; 2002)

Plerou et al. (1999; 2000a; 2000b; 2001; 2002) examined two datasets of US stocks. The first consisted of 30 minute returns for 1000 US stocks with large market capitalisation, for the two year period from 1994 to 1995, with a further two years, from 1996 to 1997, for a subset of 881 stocks. The second database involved daily returns for 422 US stocks over 35 years, from 1962 to 1996. Similar results were found for these two stock databases. After analysing the return correlation matrices with RMT, focusing mainly on the 30 minute returns, Plerou et al. (1999, 2000a,b, 2001, 2002) reached the following conclusions.

- The largest eigenvalue was found to be 25 times the maximum predicted by RMT.
- The majority (98%) of the eigenvalues were found to be within the RMT bounds, and these, along with their corresponding eigenvectors, showed good agreement with RMT.
- Eigenvectors, corresponding to the eigenvalues outside the RMT bounds, were found to display non-randomness. In fact, the authors concluded that these eigenvectors contained meaningful market information. The eigenvector for the largest eigenvalue was found to represent a “market wide” influence, in the sense that all stocks were found to participate, while almost all components had the same sign. The eigenvector for the second largest eigenvalue contained stocks with large market capitalisation relative to the database, the stocks in the third were from the electronics and computer industries, the fourth and fifth eigenvectors contained stocks involved in gold mining and investment. The sixth dealt with stocks linked to Latin America, the seventh banking firms, the eighth oil and gas, the ninth auto-manufacturing, the tenth drug-manufacturing, and the eleventh paper-manufacturing. Some of the smallest eigenvectors were found to contain single pairs of stocks with high correlations.

- The non-random eigenvectors were also found to be more stable in time, sometimes for up to 30 years. The vector for the largest eigenvalue was the most stable, while this stability decreased as the eigenvalues approached the RMT upper bound (defined Section 3.2.3, Equation 3.3).

Similar to Laloux et al. (1999, 2000), these authors also studied the application of an RMT-based filter. The filtering method was slightly different, involving replacement of the eigenvalues less than or equal to the maximum noisy eigenvalue by zeros, and preservation of the trace when restoring the main diagonal to its original values after filtering. Plerou et al. (1999, 2000a,b, 2001, 2002) followed a similar analysis to Laloux et al. (1999, 2000), splitting the data into “past” and “future” periods of equal length. The efficient frontier was then calculated with and without filtering, and the forecasted and realised risks were compared. The authors found that

- Portfolios based on the unfiltered correlation forecasts underestimated the realised risk by a factor of 2.7.
- After filtering, the forecasts underestimated by a factor of 1.25.
- The realised risks of the filtered portfolios were below those for the unfiltered portfolios.

2.4.3 Agreement between Laloux et al. (1999; 2000) and Plerou et al. (1999; 2000a; 2000b; 2001; 2002)

Working simultaneously on applying RMT to finance, these groups agreed on a number of points, central to this thesis, as follows.

- Having analysed US stock markets, the eigenvalues of the correlation matrix of returns, were found to be consistent with those calculated using random returns, with the exception of a small percentage of eigenvalues.

- The large eigenvalues, which did not conform to random returns, had eigenvectors that were more stable over time².
- Of particular interest: filtering techniques, based on RMT, were demonstrated to be beneficial in portfolio optimisation, both reducing the realised risk of optimised portfolios, and improving the forecast of this realised risk.

2.4.4 Other Work on RMT for Large Stock Portfolios

Since these key efforts, similar results have also been found for other markets. For example, Utsugi et al. (2004) investigated the Tokyo stock exchange (TSE) and the conclusions of Laloux et al. (1999) and Plerou et al. (1999, 2002) were also found to apply to this market. Using daily returns for 493 TSE stocks over a time period of 1848 days between January 1993 and June 2001, most of the eigenvalues were found by Utsugi et al. (2004) to agree with RMT and large deviating eigenvalues (and eigenvectors) were also identified.

Emerging markets have also been studied. Nilantha et al. (2007) found agreement between RMT predictions and stocks on the Sri Lankan stock exchange. They used daily returns of both the All Share and Milanka price indices, from August 2004 to March 2005, with 150 stocks considered in both cases. Again, agreement was found between RMT and most of the eigenvalues, and their eigenvectors, while large deviating eigenvalues also occurred.

Wilcox and Gebbie (2007) looked at the Johannesburg stock exchange (JSE) between January 1993 and December 2002. In this case, they had to overcome difficulties with missing data, and illiquid stocks³. In the presence of such data, good agreement with RMT was still found. For all three data cleaning methods tested, the authors reported that most eigenvalues agreed with RMT, while there were a small number of large deviating values. They also noted that the choice of data cleaning method affected the results, with some introducing more noise than others. The portfolios considered here were also large, containing

²and thus forecasts made here with matrices which have been filtered using RMT (i.e. by keeping the stable eigenvalues and smoothing the unstable ones) are expected to show greater reliability over time.

³i.e. stocks that traded irregularly

between 250 and 350 shares.

2.4.5 Work of Pafka and Kondor (2002a,b) and Pafka et al. (2004)

The work of Pafka and Kondor (2002a,b) and Pafka et al. (2004) is of relevance to us, since it introduced the concepts for filtering exponentially weighted matrices using RMT. In Pafka and Kondor (2002a,b), the authors first addressed the apparent contradiction between the widespread use of covariance matrices in finance, and the conclusions of Laloux et al. (1999; 2000) and Plerou et al. (1999; 2000a; 2000b; 2001; 2002) that use of correlation matrices which contain such a high level of noise was dangerous. Subsequently, in Pafka et al. (2004), the theory of RMT was extended to include exponentially weighted Riskmetrics style matrices. The effect, of RMT filtering on the realised risk of optimal portfolios, was also assessed there.

2.4.5.1 Pafka and Kondor (2002a,b)

In Pafka and Kondor (2002a,b), the effect of noise on portfolio risk was studied using simulations. Noiseless covariance matrices were first constructed, and then noise was added to these. Thus the true (noiseless) risk was known in advance, and so the effect of noise could be calculated, through the risk of minimum risk portfolios. The minimum risk portfolio problem aims to minimise risk over all levels of return. The advantage of using it here is that it does not require, as inputs, the forecasted returns, which contain a high level of noise themselves, and so additional noise is not introduced into the system.

One of the conclusions of this work was that noise has different effects depending on whether the risk of a portfolio is being measured, or a portfolio is being optimised. In the latter case, noise has an important effect, even for large ratios of data to assets, while in the former case it becomes insignificant a lot sooner. When portfolios are optimised in the presence of noise, the weightings of the optimal portfolio can be quite different from the true optimum. Despite this, the impact of noise on the *risk* of the optimised portfolios was concluded to be of second order. Since realised risk is the ultimate measure of investor

satisfaction in this context, the authors conclusion was that *in practice*, the risk of optimal portfolios, calculated in the presence of noise, is satisfactory.

It was shown that the measured effect of noise could be reduced in these simulations by increasing T/N , the ratio of data to assets. For a portfolio with $N = 500$ assets, a range of $2000 < T < 5000$ was considered practical, and it was concluded that in this range, the effect on the risk of optimal portfolios was of the order of 5% to 15%. The large impact of noise measured in Laloux et al. (1999, 2000) and Plerou et al. (1999, 2002, 2000a,b, 2001) was thus concluded to be the result of having a very low value of T/N in these studies. As an aside, the authors also concluded that the realised risk of these portfolios was a very good proxy for the true risk, and could be used when the true covariance matrix was not known.

While these simulation results of Pafka and Kondor (2002a,b) were indicative, there were a number of aspects which rendered the approach less than satisfactory for assessing the realised risk of medium term investments. For example, the covariance matrices constructed were simplified, and cannot therefore account for the full market structure. Also, since the market is heteroskedastic, it cannot be assumed that past data, even if available in sufficient quantities as in this study, is in fact relevant. This is compounded by the fact that the scale of forecasting period, we were most interested in, is much shorter than those considered in these simulations. The simulations also only considered the effect of noise at each level of $Q = T/N$. They did not take into consideration the possibility that the optimal value of Q may be quite different in the presence of noise than without, (which we will later find can be the case, and can have a large impact). Finally, all the RMT work previously discussed only considered the case of equally weighted covariance. Many of these questions were discussed in the subsequent work of Pafka et al. (2004).

2.4.5.2 Pafka et al. (2004)

In this later work, Pafka et al. (2004) extended RMT to Riskmetrics (1996) type financial forecasts. They derived a method for estimating the spectrum of the eigenvalues of an

exponentially weighted random matrix. The spectrum is valid in a limiting case⁴, as the number of assets becomes infinite, and the decay factor tends toward unity. Good agreement was shown between that limiting case and one with a finite number of assets and a realistic decay factor.

From this, it becomes possible to determine, numerically, the maximum eigenvalue of an exponentially weighted random matrix. Using this maximum random eigenvalue, RMT filters were developed, analogous to the equally weighted case. So, when filtering, the market eigenvalues above the maximum random eigenvalue were retained, and the rest filtered. By filtering financial covariance matrices in this way it was shown that RMT filters improved the realised risk of minimum risk portfolios, generated using exponentially weighted covariance forecasts. The main forecasting period considered here was 20 days, which is the main period we have used in our own tests.

The test methodology used was a bootstrapping technique. For a fixed number of assets, the assets in each portfolio and a test date were randomly sampled, with replacement. On each sampled date, equally and exponentially weighted covariance forecasts were generated, and filtered. Using these, a minimum risk portfolio was constructed in each case. Then, the average, *over many test dates*, of the realised risk of these minimum risk portfolios was calculated, and the performance of the different weighting and filtering combinations was compared. This method is in contrast to both Laloux et al. (1999, 2000) and Plerou et al. (1999, 2002, 2000a,b, 2001) for example, where a single test date in the middle of the data was considered. Pafka et al. (2004) tested six methods for forecasting risk, based on equally and exponentially weighted forecasts. The unfiltered forecasts were compared to those filtered using RMT, and those filtered by retaining only the largest eigenvalue. The results were as follows.

RMT filtered exponentially weighted forecasts resulted in the lowest risk from the six methods. The unfiltered forecasts were found to be more competitive when a low number of assets were used, or alternatively when large numbers of past data points were used (for

⁴This is also true of the corresponding equally weighted results

equal weights) or when the decay factor approached one (for exponential). The largest eigenvalue filter was the least successful method for both weightings.

In agreement with the previous simulations (Pafka and Kondor, 2002a,b), the effect of filtering was found to be reduced, for fixed N , as the value of T increased. A similar effect was noted for exponential weights, where increasing the decay factor was broadly found to reduce the effect of filtering. (Increasing the decay factor is equivalent to using more data, i.e. increasing T .) In the presence of filtering however, the best parameter values reduced the amount of data used, compared to the unfiltered case.

In general, and assuming the existence of sufficient data, the improvements seen after filtering were of the same order as, and below, the simulated effect of noise (Pafka and Kondor, 2002a,b). These authors also found that the decay factors, which produced the least risky portfolios, were higher than the range suggested by Riskmetrics and further concluded that the unfiltered Riskmetrics-recommended forecasts were unsuitable for their portfolio optimisation problem, more than doubling portfolio risk compared to the best filtered result.

2.4.6 Small Portfolios

While RMT has been extensively applied to large portfolios, like the S&P 500, some authors such as Conlon et al. (2007), have examined its application to smaller portfolios. Conlon et al. (2007) studied a portfolio consisting of 49 hedge funds, with limited historical data. This is relevant to our foreign exchange analysis, due to the limited number of assets in the hedge fund portfolio, since foreign exchange also has a limited number of tradeable assets. Unlike foreign exchange however, hedge fund analysis suffers from limited availability of data, since such funds typically only release results once a month. In the analysis of Conlon et al. (2007), 105 months of returns data were available for 49 hedge funds. In fact, due to the lack of data, it was necessary to select a subset of funds from the main database. Those funds with the longest track records were chosen. This is similar to what happens in practice, where typically a subset of funds with sufficiently long track records (usually five years or more) are selected for investment analysis.

There were a number of key similarities between the work of Conlon et al. (2007) and the stock market analysis of Plerou et al. (1999; 2000a; 2000b; 2001; 2002) and Laloux et al. (1999; 2000), for example. RMT was applied to the correlation matrix of hedge fund returns, and it was found that 6.1% of the eigenvalues were outside the RMT bounds. Moreover, the corresponding eigenvectors were found to contain meaningful groups of funds, for example one was related to currency funds. This is similar to the stock market analysis of Plerou et al. (1999; 2000a; 2000b; 2001; 2002) where eigenvectors were observed which related to specific sectors and regions, for example. The non-deviating hedge fund eigenvectors contained no such prevalent strategies, as was the case for stock market analysis.

When RMT was applied to the optimisation of hedge fund portfolios, the correlation matrix was cleaned using the method of Laloux et al. (1999; 2000). Here, the method of splitting the data into a past and future period of equal length was again used. Correlation forecasts were made using the past period and the realised risk in the future period was measured, as in Laloux et al. (2000). In this case an extra constraint was needed to prevent short selling⁵, which is not common in hedge fund of fund management.

Following application of the RMT filter, a 35% improvement was found, between the predicted and realised risk along the efficient frontier and, moreover, the realised risk for filtered forecasts was *always below* that for the unfiltered, with a consistent difference between them along the efficient frontier. Thus, the results of Conlon et al. (2007) show that even with a limited number of assets, RMT filtering can be effective in reducing realised risk.

2.5 Stability Filtering

Stability based filtering was first introduced in a recent paper by Sharifi et al. (2004). Here, covariance matrices calculated using equally weighted, high frequency (30 minute) S&P

⁵Short selling involves borrowing an asset and selling it to a third party. The asset is then bought at a later date and returned to the lender. The trade will profit as the price of the borrowed asset declines. Note that this is not relevant to foreign exchange, where currencies are traded in pairs, and so a long position in one currency requires a short position in another, by default.

500 returns were considered. In agreement with Laloux et al. (1999; 2000) and Plerou et al. (1999; 2000a; 2000b; 2001; 2002), it was reported that almost all of the eigenvalues of the correlation matrix conformed to those from a corresponding random matrix. In this case, 4.7% of the eigenvalues were found to lie beyond the RMT bounds.

In this work, Sharifi et al. (2004) proposed an alternative eigenvalue-filtering method, based on a principal components technique developed by Krzanowski (1984) for measuring the stability of eigenvectors, in relation to small perturbations in the corresponding eigenvalues. This was in contrast to the stability over time, as measured by Laloux et al. (1999; 2000) and Plerou et al. (1999; 2000a; 2000b; 2001; 2002). Sharifi et al. (2004) concluded that filtering correlation matrices according to the method outlined in Laloux et al. (2000) reduced Krzanowski stability. This is a direct result of the definition of Krzanowski stability, since it depends on the separation of adjoining eigenvalues. Such stability is adversely affected when the eigenvalues are close together, as is the case with the filter of Laloux et al. (2000), where the eigenvalues within the RMT bounds are set to a constant value by the filter. Thus the filter proposed in Sharifi et al. (2004) aimed to maximise this stability, by adjusting the filtered eigenvalues to be equally and maximally spaced, while preserving the trace.

In this thesis we have extended the filtering technique of Sharifi et al. (2004), making the separation of the eigenvalues an adjustable parameter. During our tests, we found that maximising Krzanowski stability led to much reduced optimisation performance, while some filters with reduced stability performed well in terms of reducing risk. By adjusting the stability of the filtered matrices, we achieved a balance between stability and risk, which improved on the stability of the commonly used filters of Laloux et al. (1999; 2000) and Plerou et al. (1999; 2000a; 2000b; 2001; 2002), while at the same time improving on the risk profile of the filter of Sharifi et al. (2004). In many of the tested cases, this extended filter performed best overall in key scenarios, including minimising realised risk, frequency of risk reduction and lowest number of large daily risk increases.

2.6 Eigenvalue Analysis

Eigenvalue analysis has also yielded other interesting results about market behaviour. In Keogh et al. (2003) analysis of the largest eigenvalue, of the covariance of Dow Jones EURO STOXX sector data, was found to reveal epochs in the market evolution. Changes in this eigenvalue were found to display a linear relationship with price within epochs. These epochs were more strongly defined for certain sectors, such as technology and telecomms.

Kwapień et al. (2004) have studied eigenvalue behaviour at different timescales for both US and German stocks. Studying timescales from seconds up to two days, they concluded that the magnitude of the leading eigenvalues increased significantly with increasing time scale. Moreover, they found that significant correlations have arisen at much smaller time scales in more recent data. The scale that contained the most information in the leading eigenvalue (days) is the one we have studied in this thesis.

In Sharkasi et al. (2006a) the two leading covariance eigenvalues, and their ratio, were examined. This revealed different responses to crashes, between mature and emerging markets. Mature markets were found to move together after a crash, and were also seen to recover faster. This analysis was further developed in Sharkasi et al. (2006b), where the authors considered the three leading eigenvalues, and their ratios. In this later work the authors also examined the eigenvalue behaviour at different timescales, using wavelet analysis. Overall, this series of work showed that the second and third largest eigenvalues held meaningful information about market events, and that this was especially true for emerging markets.

In light of these previous studies, it is the leading eigenvalue(s) that have been found to reveal information about the underlying markets. In this thesis, we review the applicability of three random matrix filters, which have been designed to isolate the information in the leading eigenvalues, i.e. those which lie beyond the range of those for random matrices.

2.7 Summary

In this chapter we have reviewed the key literature related to this thesis. We have seen that covariances vary through time, and that their estimation suffers from noise. Furthermore, this noise is expected to affect optimised portfolios, by increasing realised risk. We have reviewed two methods of forecasting covariance matrices, equally and exponentially weighted. Both of these can be filtered using concepts from RMT. Here, most eigenvalues are filtered, since they correspond well with the statistics of a random matrix. The largest eigenvalues are maintained during this filtering since they, along with their eigenvectors, have been found to contain valuable information. While much of this analysis has been performed on large systems, such as those encountered in the stock market, we have seen that filtering can also be successful when there are less assets.

We have reviewed three RMT filters, the latest of which has been based on the stability of the filtered matrix. These filters were

Laloux et al. (2000) which filters the noisy eigenvalues by replacing them with their average

Plerou et al. (2002) which filters by replacing the noisy eigenvalues with zeros, followed by restoration of the original main diagonal

Sharifi et al. (2004) which replaces the noisy eigenvalues with ones which are maximally and equally spaced and which maintain the sum of the eigenvalues. In Section 3.2.6.2, we extend this filter to vary the spacing of the replacement eigenvalues, for improved risk reduction.

In the next chapter we will review in more detail the mathematical background of this work.

Chapter 3

Background and Methodologies

3.1 Introduction

In this chapter we review key methods, used in the thesis. We first outline some financial context, discussing hedge fund investing and the foreign exchange market. This is followed by a discussion of RMT. We review RMT as applied to equally weighted matrices, and then exponential. We then discuss the RMT filters of Laloux et al. (2000), Plerou et al. (2002), and Sharifi et al. (2004) and define Krzanowski (1984) stability. We also define a novel stability-based filter. Finally, we discuss the in-sample and out-of-sample tests used to examine these filters, and the reasons these tests were chosen.

3.2 Background

3.2.1 Hedge Fund Portfolio Investment

A hedge fund (e.g. L'Habitant (2006)) is a private investment vehicle typically only available to accredited investors, namely high net worth individuals and institutional investors (e.g. commercial banks and pension funds). This pool of investors are able to seek more sophisticated strategies, requiring a deeper understanding on the part of the investor, than are typically available to the general public. Hedge funds come in many varieties, involving

a wide range of investment strategies, but one common technique involves the forecasting of asset returns, and their associated risks and correlations, and the construction of optimal investment portfolios, based on these forecasts (e.g. Elton et al. (2006)).

In this work, we have focused on improving the risk profile of two common investment scenarios. The first is the buying and short selling of stocks, in this case US stocks from the S&P 500. The second scenario is the buying and selling of currencies and commodities. For this research, we were principally concerned with managed, medium term investments in these markets. We have endeavoured to mimic the active investment strategies of hedge funds with our research methods, leading to our use of forward validation. At the same time, we have been careful to avoid strategy specific assumptions, instead preferring a test environment which has relevance to potentially many different investment strategies. This aspect is discussed further in Section 3.3. Finally, while hedge fund style investing has been the main focus of this work, these results are also relevant to other market participants, if employed correctly.

3.2.2 Foreign Exchange Market

While most of the literature on RMT is concentrated on stock market analysis, as seen in Chapter 2, in this work we have taken filtering methods developed in the stock market, and applied them in the foreign exchange (“Fx”) market. We found remarkable agreement between the effects of filtering in these markets.

The Fx market is by far the biggest and most liquid market in the world. According to the Bank for International Settlements Triennial Report (Heath et al., 2007) the average daily traditional¹ turnover grew by 69% to to \$3.21 trillion over the previous three years. This was an unprecedented rate of growth, much larger than the previous period. In this time, the type of market activity seen was found to have changed substantially. Transactions with financial institutions like hedge funds, mutual funds, pension funds and insurance companies more than doubled, while more diversification of the currencies being traded

¹Traditional turnover consists of spot trades, outright forwards, and fx swaps.

was observed. It is in this climate that we study the application of RMT filters to foreign exchange.

3.2.3 RMT and Historical Covariance

In this section we review the statistical behaviour of the eigenvalues of a random matrix constructed using equally weighted returns. (In particular, we require the expression for the *extremal eigenvalue*, λ_+ , of these equally weighted random matrices, in the filtering algorithms that follow.) These statistics have been found to correspond closely with the bulk of observed market eigenvalues, as discussed in Chapter 2.

As described by Laloux et al. (1999), Plerou et al. (2002), Sharifi et al. (2004) and others, in the context of correlation matrices of financial returns, if \mathbf{R} is any matrix defined by

$$\mathbf{R} = \frac{1}{T} \mathbf{A} \mathbf{A}' \quad (3.1)$$

where \mathbf{A} is an $N \times T$ matrix whose elements are i.i.d.² random variables with zero mean and finite variance, then it has been shown (Sengupta and Mitra, 1999) that, in the limit $N \rightarrow \infty, T \rightarrow \infty$ such that $Q = T/N \geq 1$ is fixed, the probability density function $P(\lambda)$ of the eigenvalues of \mathbf{R} is self-averaging³, and is given by

$$P(\lambda) = \begin{cases} \frac{Q}{2\pi\sigma^2} \frac{\sqrt{(\lambda_+ - \lambda)(\lambda - \lambda_-)}}{\lambda} & \text{if } \lambda_- \leq \lambda \leq \lambda_+ \\ 0 & \text{otherwise} \end{cases} \quad (3.2)$$

where σ^2 is the variance of the elements of \mathbf{A} , and the extremal eigenvalues are given by

$$\lambda_{\pm} = \sigma^2 \left(1 + 1/Q \pm 2\sqrt{1/Q} \right). \quad (3.3)$$

²i.i.d. \equiv independent and identically distributed

³In this context, the term *self-averaging* is used to indicate that, in the large limit of N and T the p.d.f of the eigenvalues of R tends to $P(\lambda)$ (as given in Equation 3.2)

Financial correlation and covariance matrices can be expressed, in general, in the form given by Equation (3.1)⁴, so matrices for historical data can be compared to those generated from i.i.d. random returns. (The details of the method for doing this are discussed in Section 3.3.1)

Here we define the *covariance* matrix $\mathbf{V} = \{\sigma_{ij}\}_{i,j=1}^N$ of returns^{5 6} by

$$\sigma_{ij} = \overline{G_{it}G_{jt}} - \overline{G_{it}}\overline{G_{jt}} \quad (3.4)$$

where $\overline{G_{it}}$ refers to the mean over time, and the *correlation* matrix $\mathbf{C} = \{\rho_{ij}\}_{i,j=1}^N$ by

$$\rho_{ij} = \frac{\sigma_{ij}}{\sqrt{\sigma_{ii}\sigma_{jj}}} \quad (3.5)$$

where $\{G_{it}\}_{t=1,\dots,T}^{i=1,\dots,N}$ are the returns

$$G_{it} = \ln \left(\frac{S_{it}}{S_{i,t-1}} \right) \quad (3.6)$$

and where S_{it} is the spot price of asset i at time t .

Equation (3.3) now provides us with a closed form for the largest eigenvalue of equally weighted random matrices, to be used with the filtering techniques.

3.2.4 RMT and Exponentially Weighted Covariance

In this section we review the behaviour of the eigenvalues of a random matrix, which has been constructed using exponential weights. Again, for the filtering algorithms that follow we will need to calculate the largest eigenvalues of such matrices.

In extending RMT filtering to exponentially weighted matrices, Pafka et al. (2004) have

⁴where N is the number of assets and T is the number of past price moves (one move is recorded each day in these markets)

⁵throughout this thesis the following notation is used: $\{x_i\}_{i=1}^N \equiv \{x_i : i = 1, \dots, N\}$, $\{x_{ij}\}_{i,j=1}^N \equiv \{x_{ij} : i = 1, \dots, N; j = 1, \dots, N\}$, $\{x_{it}\}_{t=1,\dots,T}^{i=1,\dots,N} \equiv \{x_{it} : i = 1, \dots, N; t = 1, \dots, T\}$ etc.

⁶Please note that, formally, objects such as σ_{ij} and ρ_{ij} can be considered as having a time subscript, as they will vary through time by definition. However, this subscript has been dropped in what follows, so as not to clutter the already complex notation.

analysed matrices of the form $\mathbf{M} = \{m_{ij}\}_{i,j=1}^N$ with

$$m_{ij} = \sum_{k=0}^{\infty} (1 - \alpha) \alpha^k x_{ik} x_{jk} \quad (3.7)$$

and where $\{x_{ik}\}_{k=0,\dots,\infty}^{i=1,\dots,N}$ are assumed to be $N.I.D.(0, \sigma^2)$ ⁷. They have shown that, in the special case $N \rightarrow \infty$, $\alpha \rightarrow 1$ with $Q \equiv 1/(N(1 - \alpha))$ fixed, the probability density function, $\rho(\lambda)$, of the eigenvalues of \mathbf{M} is given by

$$\rho(\lambda) = \frac{Qv}{\pi} \quad (3.8)$$

where v is the root of

$$F(v) = \frac{\lambda}{\sigma^2} - \frac{v\lambda}{\tan(v\lambda)} + \ln(v\sigma^2) - \ln(\sin(v\lambda)) - \frac{1}{Q} \quad (3.9)$$

$F(v)$ is well defined on the open interval $(0, \pi/\lambda)$. If a root does not exist on this interval for a given value of λ we define $\rho(\lambda) = 0$ for that λ . The family of matrices, defined by Equation (3.7), includes the Riskmetrics (1996) covariance and correlation matrices. We define this exponentially weighted covariance matrix $\mathbf{V}^* = \{\sigma_{ij}^*\}_{i,j=1}^N$ by

$$\sigma_{ij}^* = \frac{1 - \alpha}{1 - \alpha^T} \sum_{t=0}^{T-1} \alpha^t (G_{i,T-t} - \overline{G_{it}}) (G_{j,T-t} - \overline{G_{jt}}) \quad (3.10)$$

and define the corresponding, exponentially weighted, correlation matrix $\mathbf{C}^* = \{\rho_{ij}^*\}_{i,j=1}^N$ by

$$\rho_{ij}^* = \frac{\sigma_{ij}^*}{\sqrt{\sigma_{ii}^* \sigma_{jj}^*}} \quad (3.11)$$

Here, α is commonly called the *decay factor*.

The largest eigenvalue of an exponentially weighted random matrix can now be determined by examining Equation (3.9) numerically. We have derived a more efficient method (Daly et al., 2008), presented in Appendix A, where we have shown that this ex-

⁷ $N.I.D.(\mu, \sigma^2) \equiv$ Normally and identically distributed (with mean μ and variance σ^2)

tremal eigenvalue is also the solution of

$$\frac{\lambda}{\sigma^2} - \ln \left(\frac{\lambda}{\sigma^2} \right) = 1 + \frac{1}{Q}, \quad \lambda > \sigma^2 \quad (3.12)$$

3.2.5 Standard Filtering Methods

The three filtering methods compared in this thesis are all based on replacing the “noisy” eigenvalues of the correlation matrix, while maintaining its trace. The noisy eigenvalues are taken to be those that are less than the largest eigenvalue of the corresponding random matrix⁸. This is because, as discussed in Chapter 2, the statistics of those eigenvalues below λ_+ have been observed to correspond closely with RMT, for different markets. The empirical conclusion has been that only the largest market eigenvalues (those above λ_+) can be relied upon to contain genuine information (c.f. Chapter 2).

There are a number of empirical reasons why the minimum random eigenvalue is not considered in these filters. First amongst them is that, while the large market eigenvalues which have been studied are *clearly* separated from the RMT bounds, the same cannot be said for the smallest measured eigenvalues. Typically, small eigenvalues have been found outside the bounds, which were consistent with those which occur simply due to the fact that N and T are finite in practice. To complement this, following examination of the *eigenvectors* corresponding to the large eigenvalues, *clear* non-randomness, and stability over time, of the eigenvectors has been reported, while the same cannot be said for the measured eigenvectors of the smallest eigenvalues. For these reasons, we follow the suggestions of the original authors in this field (Plerou et al. (2002), Pafka et al. (2004), Laloux et al. (2000)) and retain only the eigenvalues above λ_+ during filtering. We suggest that further examination of the minimum RMT eigenvalue would only be valid, and potentially useful, if it were done in conjunction with the future work suggested in Section 6.3, to further study the effect of finite dimensions (N and T).

⁸The corresponding random matrix is the random matrix which uses equivalent values of Q and σ^2 . This is further discussed in Section 3.3.1

The theoretical limiting cases (described in Sections 3.2.3 and 3.2.4) are commonly used to estimate the largest eigenvalues of the random matrices, which can also be estimated by calculating them directly from Monte Carlo simulated random returns. In this work, we have used the maximum eigenvalue predicted by RMT (Equations 3.3 and 3.12) for filtering, due to computational efficiency. We now discuss the three RMT filters which use this concept.

3.2.5.1 LCPB Filtering Method

The filtering method of Laloux et al. (2000) (and referred to here as LCPB) replaces the noisy eigenvalues with their mean as follows. Starting with the sequence, $\Lambda = \{y_i\}_{i=1}^N$, of eigenvalues (ordered by size of the numbers), of some $N \times N$ matrix, \mathbf{M} , and the corresponding eigenvectors, \mathbf{E} , we define the sub-sequence

$$\Lambda_{\text{noisy}} = \{y_i\}_{i=1}^n \quad (3.13)$$

of eigenvalues which are less than the maximum eigenvalue predicted by RMT. A sequence of filtered eigenvalues are then defined as

$$\Lambda_{\text{filtered}} = \{x_1, \dots, x_n, y_{n+1}, \dots, y_N\} \quad (3.14)$$

where $\{y_i\}_{i=n+1}^N$ are the eigenvalues assumed to contain information and

$$\Lambda_{\text{new}} = \{x_i\}_{i=1}^n \quad (3.15)$$

are the replacements, where for all $i = 1, \dots, n$ we have

$$x_i = \frac{1}{n} \sum_{j=1}^n y_j \quad (3.16)$$

These filtered eigenvalues $\Lambda_{\text{filtered}}$ are then combined with the original eigenvectors, \mathbf{E} , using the eigen decomposition theorem ⁹, to construct a filtered matrix

$$\mathbf{M}_{\text{filtered}} = \mathbf{E}\mathbf{D}_{\text{filtered}}\mathbf{E}^{-1} \quad (3.17)$$

where $\mathbf{D}_{\text{filtered}}$ is a matrix with $\Lambda_{\text{filtered}}$ on the main diagonal and zeros everywhere else. Replacing the noisy eigenvalues by the mean noisy eigenvalue means that the trace of $\mathbf{M}_{\text{filtered}}$ is equal to the trace of \mathbf{M} . When using these three RMT filters the trace (i.e. the sum of the eigenvalues) should be preserved so that the system does not distort.

3.2.5.2 PG+ Filtering Method

As described by Plerou et al. (2002) (and referred to here as PG+), this method is the same as the LCPB method, except that the noisy eigenvalues are all replaced by zeros. Then, after the filtered matrix $\mathbf{M}_{\text{filtered}}$ is built, its main diagonal is set to be equal to that of the original matrix \mathbf{M} , thus preserving the trace (i.e. the sum of the eigenvalues) to prevent system distortion.

3.2.6 Stability Based Filtering

In this section we give the details of our novel, stability-based filter, which is an extension of that of Sharifi et al. (2004), which is in turn based on the stability (defined below), as described by Krzanowski (1984), of the filtered matrix.

3.2.6.1 Krzanowski Stability

Krzanowski (1984) measured eigenvector stability, specifically the effect on each eigenvector of a perturbation in the corresponding eigenvalue. This is in contrast to stability over time, as analysed by many other authors, e.g. Laloux et al. (2000) and Plerou et al. (2002).

⁹Let \mathbf{M} be a square matrix and let \mathbf{E} be a matrix of eigenvectors. If \mathbf{E} is a square matrix then $\mathbf{M} = \mathbf{E}\mathbf{D}\mathbf{E}^{-1}$ where \mathbf{D} is a diagonal matrix containing the corresponding eigenvalues on the main diagonal, (e.g. Strang, 1980).

Krzanowski (1984) considered the angle, θ_i , between an eigenvector v_i and v_i^p , where v_i^p is the maximum perturbation that can be applied to v_i while ensuring that the eigenvalue, λ_i^p , corresponding to v_i^p is within ϵ of the eigenvalue, λ_i , corresponding to v_i . He showed that θ_i is given by:

$$\cos \theta_i = \begin{cases} \left(1 + \frac{\epsilon}{\lambda_i - \lambda_{i-1}}\right)^{-\frac{1}{2}} & \text{for } \lambda_i^p < \lambda_i \\ \left(1 + \frac{\epsilon}{\lambda_{i+1} - \lambda_i}\right)^{-\frac{1}{2}} & \text{for } \lambda_i < \lambda_i^p \end{cases} \quad (3.18)$$

where $\lambda_1 \leq \lambda_2 \leq \dots \leq \lambda_L$ are the eigenvalues, and suggested using $\epsilon = k\lambda_i$ with $k = 0.1, 0.05$ or 0.01 . When measuring stability, we have chosen $k = 0.1$, which was the most consistent with typical eigenvalue changes between different subperiods of our data. Note however that the choice of k has no effect on the filtering algorithms. Rather, it is used when *measuring* the stability (Eqn 3.18) of the subsequently filtered matrices, as is done in Section 4.4.3. When measuring this mean stability, of the filtered and unfiltered covariance matrices, the arithmetic mean of the cases $\lambda_i^p < \lambda_i$ and $\lambda_i < \lambda_i^p$ was calculated.

To summarise, we consider here the stability of the eigenvectors, in response to changes in the corresponding eigenvalues. In the Krzanowski method, stability is measured by considering the biggest perturbation to an eigenvector that can be generated by a perturbation, no larger than ϵ , in the corresponding eigenvalue. In the in-sample and out-of-sample tests that follow we assess empirically how this stability impacts the performance of optimal portfolios, (i.e. calculated using underlying matrices of varying stability). We find that increasing Krzanowski stability can improve the realised risk of those optimal portfolios. It is possible that one reason for this may be that changes in the market (caused by the arrival of new information) result in eigenvalue fluctuations over the life of the investment, and so matrices which are well balanced in anticipation of such changes perform better, although this has not been proved.

3.2.6.2 Stability-based KR Filter

The stability-based filtering method is defined as follows. To maximise the Krzanowski stability of the filtered matrix, while also maintaining its trace, the method of Sharifi et al. (2004) replaces the noisy eigenvalues with ones that are equally and maximally spaced, are positive, and have sum equal to the sum of those replaced. This formulation is a result of the observation that increased stability of an eigenvector follows directly from increased separation of its eigenvalue from the neighbouring eigenvalues (c.f. Eqn 3.18). To achieve maximal spacing, the filter of Sharifi et al. (2004) assumed that the smallest replacement eigenvalue should be very close to zero.

During the work described in this thesis, the method of Sharifi et al. (2004) was adapted to create a novel filter, by making the smallest replacement eigenvalue a parameter of the filter, so that changes in stability and *optimisation results*, achieved for various values of this parameter, could be measured. We call the adapted version the *KR method*, (Daly et al., 2008). It follows that the original filter of Sharifi et al. (2004), using a minimum replacement eigenvalue of 10^{-8} , is a special case of the new KR method

The KR method is identical to the LCPB method except in the choice of eigenvalues to replace the noisy eigenvalues. If

$$\Lambda_{\text{noisy}} = \{y_i\}_{i=1}^n \quad (3.19)$$

are the original noisy eigenvalues then, for the KR method, the replacement eigenvalues

$$\Lambda_{\text{new}} = \{x_i\}_{i=1}^n \quad (3.20)$$

are given by

$$x_i = x_1 + (i - 1)k \quad (3.21)$$

for some constant $k \geq 0$, which is defined by the choice of minimum replacement eigenvalue, x_1 , and the constraint that the sum of the replacement eigenvalues must equal the

Table 3.1: List of KR methods tested

Method Name	Minimum Replacement Eigenvalue
KR2	$\frac{1}{2}\bar{\Lambda}_{\text{noisy}}$
KR4	$\frac{1}{4}\bar{\Lambda}_{\text{noisy}}$
KR8	$\frac{1}{8}\bar{\Lambda}_{\text{noisy}}$
KR16	$\frac{1}{16}\bar{\Lambda}_{\text{noisy}}$
KR64	$\frac{1}{64}\bar{\Lambda}_{\text{noisy}}$
KR100	$\frac{1}{100}\bar{\Lambda}_{\text{noisy}}$
KR1000	$\frac{1}{1000}\bar{\Lambda}_{\text{noisy}}$
KR0	10^{-8}

sum of the eigenvalues being replaced to prevent distortion. In addition, the replacement eigenvalues must all be strictly positive. It follows that

$$k = \frac{2(a - x_1)}{n - 1} \quad (3.22)$$

where a is the mean of the eigenvalues being replaced. Since $k \geq 0$ we require that $x_1 \leq a$. Moreover, the case $k = 0$ just collapses to the LCPB method, since $k = 0 \Leftrightarrow x_1 = a$.

The KR methods considered here, and their defining minimum replacement eigenvalues, are listed in Table 3.1, where $\bar{\Lambda}_{\text{noisy}}$ refers to the mean of the noisy eigenvalues. In the in-sample tests, these variations of the KR filter were compared. In the out-of-sample tests, one had to identify, based on past performance, the best KR filter to be used each day for that days investment, (further details provided in Section 3.3.2).

3.3 Methodologies

In this thesis we have employed two main methods for testing these filters. The first, in-sample, method follows directly from Pafka et al. (2004), and involves randomly sampling many test dates with replacement. The second, out-of-sample, method known as forward validation, evaluates all dates once, by simulating an actual implementation over time. We give the details of these two methods now.

3.3.1 In-sample Methodology

For any in-sample analysis, and following Pafka et al. (2004), bootstrapped samples were taken, together with the mean across these samples. For a given value of N , (the number of assets), we randomly selected N assets from the data set, and a random test date. Everything up to and including the test date was taken as historical information and everything afterward as realised, future information. For each N , we repeated this random selection 1000 times, with replacement, and calculated the mean, across all bootstrapped samples, of the realised risk of the forecast minimum risk portfolio (Pafka et al., 2004), calculated using our forecast covariance.

A covariance forecast in this context consisted of a raw forecast, which was either exponentially or equally weighted, and could be unfiltered, or filtered by one of the LCPB, PG+ or KR methods applied to the correlation or covariance matrix.

On each test date, we calculated the forecast minimum risk portfolio, optimised as follows (Pafka et al., 2004). Choose a portfolio weighting $\{w_i\}_{i=1}^N$ that minimises the total expected risk

$$\sum_{i,j=1}^N w_i w_j \hat{\sigma}_{ij} \quad (3.23)$$

while satisfying the budget constraint

$$\sum_{i=1}^N w_i = 1 \quad (3.24)$$

Here, $\hat{\mathbf{V}} = \{\hat{\sigma}_{ij}\}_{i,j=1}^N$ is the forecast covariance matrix, which can be either equally (Eqn 3.1) or exponentially (Eqn 3.7) weighted, and filtered or unfiltered. The solution, $\{\hat{w}_i\}_{i=1}^N$, of this problem is:

$$\hat{w}_i = \frac{\sum_{j=1}^N \hat{\sigma}_{ij}^{-1}}{\sum_{j,k=1}^N \hat{\sigma}_{jk}^{-1}} \quad \forall i \quad (3.25)$$

where $\hat{\mathbf{V}}^{-1} = \{\hat{\sigma}_{ij}^{-1}\}_{i,j=1}^N$ is the matrix inverse of $\hat{\mathbf{V}}$. The *realised risk* of the optimal portfolio is then defined by

$$\sqrt{\sum_{i,j=1}^N \hat{w}_i \hat{w}_j \tilde{\sigma}_{ij}} \quad (3.26)$$

where $\tilde{\mathbf{V}} = \{\tilde{\sigma}_{ij}\}_{i,j=1}^N$ is the realised covariance matrix, and is the (equally weighted) covariance matrix of the realised future returns over the investment period, calculated analogously to Equation 3.4. (This realised risk thus describes the variation in returns of the investment. Overall investment success is typically measured by its return divided by this risk.) The *forecast risk* is calculated analogously, using the forecast covariance matrix, $\hat{\mathbf{V}}$. As seen in Pafka and Kondor (2002a,b), (discussed Section 2.4.5.1), this realised risk can be considered a good proxy for the true portfolio risk, which can never be known.

By comparing the covariance forecasts in this way, we measured their effect on realised risk without using forecast returns, which would introduce additional noise into the results. Further, we have not used any knowledge of future returns in our tests, since we wished to evaluate both forecasting methods (equal vs. exponential weighting) as well as filtering methods. This is in contrast to some previous studies (e.g. Laloux et al. (2000)), that have isolated the effect of the filtering method on the correlation matrix, by using future knowledge of realised returns to estimate the variance of each individual asset.

We now detail the bootstrapping procedure, referring to the mathematical methods defined earlier in this chapter.

For each portfolio size considered (e.g. $N = 100, 250$, or 432 for stocks), we:

1. First select a random test date (with replacement) from those available
2. For that date, select a random list of N assets (again with replacement) from the full list of available assets. (This avoids having to forecast returns, which would introduce further noise.)
3. For that list of assets, on that date, we generate alternative forecast covariance matrices, $\hat{\mathbf{V}}$, as follows.

- (a) We start with an unfiltered covariance matrix, which is calculated from either equally (c.f. Eqn 3.4) or exponentially (c.f. Eqn 3.10) weighted past returns.
- (b) We apply the filters by first selecting a target matrix. We have tested filtering the covariance matrix directly, as well as filtering the associated correlation matrix. The latter is more standard. To perform this filtering we first require the maximum eigenvalue of the *corresponding* random matrix. This maximum eigenvalue can be found using Eqn 3.3 for the equally weighted case, or Eqn 3.12 for exponential.

In both cases, for calculating the appropriate λ_+ , we require two inputs, namely Q and σ . The calculation of Q is trivial for both weighting schemes. We calculate σ in the equally weighted case as follows.

- (c) As defined in Section 3.2.3, σ^2 is the variance of the elements of A (where A is first described as part of Eqn 3.1). All that is required here is to transform our target matrix into the form of R specified in Eqn 3.1. For filtering correlation we set $A = \{a_{it}\}_{t=1,\dots,T}^{i=1,\dots,N}$ where

$$a_{it} = \frac{G_{it} - \overline{G_{it}}}{\sqrt{\sigma_{ii}}} \quad (3.27)$$

and for filtering covariance we have $A = \{a_{it}\}_{t=1,\dots,T}^{i=1,\dots,N}$ where

$$a_{it} = G_{it} - \overline{G_{it}} \quad (3.28)$$

where σ_{ii} and G_{it} are defined in equations 3.4 and 3.6 respectively, and where we recall that $\overline{G_{it}}$ refers to the mean over time. σ^2 is now the variance of all the elements of A combined. Identical transformations are required in the exponentially weighted case to calculate the appropriate σ^2 .

We can now calculate λ_+ for the target matrix using the appropriate equation (Eqn 3.3 for equally weights or Eqn 3.12 for exponential).

- (d) Once λ_+ has been calculated, we can apply the filtering techniques defined in Sections 3.2.5.1 (LCPB method), 3.2.5.2 (PG+ method) and 3.2.6.2 (KR method) to the target matrix. Note that the treatment of the KR filter in the *in-sample* (bootstrapping) case involves considering *separately* all possible flavours of the KR filter (c.f. Section 3.2.6.2), which have been listed in Table 3.1. In this table we see that we are considering a total of eight different KR filters here (KR2, KR4, etc), and we compare them to each other in-sample, and also to the LCPB and PG+ methods.
- (e) Finally, when correlation has been the target matrix for filtering, we calculate our forecast covariance matrix $\hat{\mathbf{V}} = \{\hat{\sigma}\}_{i,j=1}^N$, as usual, using the transformation

$$\hat{\sigma}_{ij} = \hat{\rho}_{ij} \sqrt{\sigma_{ii} \sigma_{jj}} \quad (3.29)$$

where $\{\hat{\rho}\}_{i,j=1}^N$ is the result of filtering the correlation matrix using RMT (i.e. in step 3(d)), while σ_{ii} is as before, (Eqn 3.4 when using equal weights and Eqn 3.10 when using exponential).

4. Having calculated these filtered covariance matrices in this way, we can now compare them with the unfiltered covariance as follows. We evaluate each forecast covariance matrix (each day) by calculating the realised risk (c.f. Eqn 3.26) of the corresponding minimum risk portfolio (c.f. Eqn 3.25) for each forecast matrix (for that day).
5. We now repeat steps 1-4 one thousand times and calculate, for each method of forecasting covariance, the average realised risk (as calculated in step 4) across these

experiments.

6. These average risks are then examined for different values of the decay factor (α) (for exponential weights) or size of time window (T) (for equal weights), and the filters are compared directly. An example of such a comparison can be seen in Figure 4.5.

While this method yields valuable information about the average behaviour of the filters, it should be noted that, since it is an in-sample method (i.e. best parameter values are assessed *after* the test), the KR method has some advantage over both the PG+ and LCPB methods. This is because it has an extra parameter (namely its minimum replacement eigenvalue, leading to alternative specifications KR2, KR4, etc as listed in Table 3.1) and thus the KR method has more potential for “fitting the data”. Despite this, we subsequently found good agreement between the results of the in-sample tests and those of the out-of-sample tests described in the following section.

The results of such in-sample risk analysis are found in Sections 4.4.4, 4.6.1, 4.7.1, and 5.5.

3.3.2 Out-of-sample Methodology

For comparing the forecasting and filtering models out-of-sample we used forward validation. This method simulates a live implementation (of trading every day using risk forecasting). The method considers every available test date and, for each one, uses only data (and past performance) prior to the test date to optimise any model parameters, as follows. (Note that this method is very similar to the bootstrapping method of Section 3.3.1, so we will concentrate on the differences here.)

1. We first make a selection of assets to examine. For both markets we have examined the full list of available assets (results found in Sections 4.5.2, 4.6.2, and 5.6). For foreign exchange we also examined a subset of major assets (results found in Section 5.6.5). (Again, asset selection is required to avoid the noise introduced when using forecast returns.)

2. This is a preparatory stage, which sets up all the results which we need during the forward validation. For each filter (LCPB, PG+, KR, and unfiltered), and for both weighting schemes (equal and exponential), for all available test dates, and for all parameter values considered (i.e. values of T and α , with the KR filter having a second parameter, x_1 , the minimum replacement eigenvalue), we implement steps 3 and 4 of the in-sample algorithm. This means that we:

- generate all possible covariance forecasts (step 3)
- and calculate the realised risk of the minimum risk portfolio associated with each forecast (step 4).

Note that for in-sample testing we have used 1000 randomly sampled dates, while in the out-of-sample algorithm here we have assessed *each available date once*. Note also that the treatment of the KR filter has now changed from the in-sample case. Here, we are no longer considering KR2, KR4 etc. explicitly. Instead, we consider the different KR flavours as parameterisations of the more general “KR” method, with each one defined by its minimum replacement eigenvalue x_1 (c.f. Section 3.2.6.2).

To summarise: in this step, we have generated the following data set (which we will refer to in the steps that follow): $\{U_{t',f,w,p}\}$ where U refers to the realised risk (c.f. Eqn 3.26) of the minimum risk portfolio (c.f. Eqn 3.25) associated with some covariance forecast, t' refers to the test date (the start date of the investment), f refers to the choice of filter used in the forecast (LCPB, PG+, KR, or unfiltered), w refers to the choice of weighting scheme (equal or exponential) and p refers to the parameters associated with that weighting scheme and filter combination.

3. Using this, we now forward validate as follows. For a test date t' , for a given filter, f , and weighting scheme, w , we choose the optimal parameters $p'_{t'}$ which minimises *average past realised risk* associated with this forecasting method. We effectively

minimise this sum:

$$\sum_{t \leq t' - t_F} U_{t,f,w,p} \quad (3.30)$$

Note that we only consider portfolios which were invested before $t' - t_F$, or on that date, where t_F is the investment period (20 days) since, on day t' , we do not have full knowledge of investments made after $t' - t_F$, since conceptually they are still invested.

4. Now, using the optimal parameters generated in the previous step, we invest forward using these parameters and, for each $\{t', f, w\}$ we measure $U_{t',f,w,p_{t'}}$, (which of course has been pre-calculated as part of step 2). This quantity is therefore the realised portfolio risk (c.f. Eqn 3.25) of the minimum risk portfolio (c.f. Eqn 3.25) which is associated with the covariance forecasting method $\{f, w\}$ on day t' .

We study the behaviour of these realised risks by examining their average behaviour (e.g. Section 4.5.2.1), their average behaviour in each year (e.g. Section 4.5.2.2), and their individual behaviours (e.g. 4.5.2.3), as would be done in practice when considering an investment.

We now have a simulation of how each forecasting method would have performed in practice. This method allows the fair comparison of filtering methods with different numbers of parameters and also gives some insight into the stability of the models over time. In comparison, as previously stated, the bootstrapping method can favour models with more parameters, since the best parameter values are assessed *after* testing. In this work, the stability-based filter had one extra parameter, the choice of the minimum replacement eigenvalue.

Note that the forward validation test also allows for the optimal parameters to vary over time, which is realistic, since markets are capable of rapidly changing behaviours, (and so choosing one optimum parameter after testing may not be appropriate). In contrast, when performing the bootstrapping procedure, all parameter combinations were tested, and compared after the fact, and so any “best” combination is assumed to have been at fixed

values over the entire test period.

Forward validation, also known as historical simulation or back testing, is a typical test performed prior to implementing any new models into a systematic trading environment. Its primary aim is to simulate a live deployment. In our case, daily parameter retraining reflects the trading style of a highly adaptable, active fund manager.

3.4 Summary of Scenarios Tested

In Chapters 4 and 5 various test scenarios are explored, for measuring the risk reducing capacity of the RMT filters. In Tables 3.2 and 3.3, we summarise the different variations of assets and target matrices tested in these scenarios, and state where the relevant key results are reported in the thesis. Table 3.2 refers to in-sample testing and Table 3.3 to out-of-sample testing. In **all** cases we have tested both equal and exponential weighting systems.

Table 3.2: List of **in-sample** scenarios used to test the effect of filtering on risk, and where the main results are reported, for different portfolio sizes (N). Three different categories have been tested: (1) filtering the S&P correlation, (2) filtering the S&P covariance, and (3) filtering the Fx correlation.

N	S&P 500 / Corr	S&P 500 / Cov	Fx / Corr
432	Section 4.4	Section 4.6	-
250	Section 4.4	Section 4.6	-
100	Section 4.4	Section 4.6	-
39	Section 5.5	-	Section 5.5
30	-	-	Section 5.5
20	-	-	Section 5.5
15	-	-	Section 5.5
10	-	-	Section 5.5

Table 3.3: List of **out-of-sample** scenarios used to test the effect of filtering on risk, and where the main results are reported, for different portfolio sizes (N). Three different categories have been tested: (1) filtering the S&P correlation, (2) filtering the S&P covariance, and (3) filtering the Fx correlation.

N	S&P 500 / Corr	S&P 500 / Cov	Fx / Corr
432	Section 4.5	Section 4.6	-
39	-	-	Section 5.6
15	-	-	Section 5.6

3.5 Summary

This chapter has outlined the different models and techniques that have been used in this thesis to test the RMT filters. It has also shown how a novel, stability-based filter may be developed. In the forthcoming chapters, this novel filter has been compared, in-sample and out-of-sample, to the two well known filters also described in this chapter, first in the context of stock market trading, and then foreign exchange. This stability-based filter was developed to improve, compared to pre-existing methods, the stability of optimised solutions, in response to changes in the market eigenvalues.

The testing methods of this thesis also aimed to improve on the available literature, in particular by providing an out-of-sample test of the models using forward validation. The advantages of this method are that

1. forward validation simulates an actual implementation of the test models.
2. by choosing parameters each day based only on “past data” it allows the comparison of models with different numbers of parameters, and avoids “over fitting” of the data.
3. it does not assume that model parameters should be fixed over time. This is important, since it allows for evolution in the market. For example, there may be changes to prevailing economic policies.

4. it allows the generation of meaningful annual, monthly etc. averages, which would be required by practitioners. While overall performance is important, consistency is also valued.

In the next chapter we use these in-sample and out-of-sample methods to evaluate the application of RMT filtering to a stock portfolio.

Chapter 4

RMT for Large Stock Portfolios

4.1 Introduction

In this chapter we consider the case of a portfolio manager trading S&P 500 stocks, and evaluate whether RMT filtering can be used to reduce risk. Since the number of assets tested was large in comparison to the amount of available data points, we expected to be able to filter noise from our forecasts of risk. We compared the three RMT filters described in Sections 3.2.5 and 3.2.6.2 (the LCPB, PG+ and KR filters) to establish how each filter could be used to improve managed, medium term equity investments. We first recall the work of Pafka et al. (2004) in this area, which is particularly relevant to the work that follows. Pafka et al. (2004) first studied the application of these filters to medium term stock portfolios, including exponentially weighted forecasts in this analysis.

The results here are separated into in-sample and out-of-sample. For the in-sample results we started by assessing the noise in the unfiltered forecasts, and comparing the stability of the filtered matrices, for the different filters. We then used this bootstrapping technique to assess the effect of the filters on the realised portfolio risk, and we compared our results to those of Pafka et al. (2004), where we also considered stability-based filters in our comparison. Here we found that our modified stability-based filter outperformed both the unfiltered results, and the other filters tested, for three different sizes of portfolio, containing 100, 250,

and 432 assets respectively. How these stocks were chosen is documented below.

We then tested these filters using forward validation, which was our primary testing methodology. Forward validation allowed us to better compare filters with different numbers of parameters, and best replicated an actual implementation of the methods used here. The forward validation was done for the full portfolio of 432 assets. We again found that our novel, stability-based filtering produced the best average realised risks, while it also generated the highest frequency of reduction, and reduced the number of days that filtering resulted in large risk increases. These daily increases had not been previously discussed in the literature on RMT. We also analyse their effect on the range of the daily realised risk.

In comparison to the filtering tested above, we have also tested the concept of filtering the covariance matrix directly, since this matrix was found to contain more information than the correlation matrix. Here we have found that, while both correlation and covariance filtering improved risk out-of-sample (compared to the unfiltered) in all cases, correlation filtering was preferred on average.

In conjunction with these tests, we also analysed the parameter values recommended by both the bootstrapping, and forward validation, and compared these to Pafka et al. (2004) and Riskmetrics (1996), which are in disagreement on the best choice of exponential decay factor in particular.

We finish by comparing all the forecasting methods, filtered and unfiltered, discussed in this chapter.

4.2 Review of Pafka et al.

In this section we review the work of Pafka et al. (2004), and its relevance to this thesis. The authors examined six methods of forecasting risk. Starting with equally and exponentially weighted covariance forecasts, they tested standard random matrix filtering alongside a filter which retained only the largest eigenvalue, and compared these results with the unfiltered forecasts. Using the bootstrapping technique described in Section 3.3.1, Pafka et al. (2004),

analysed portfolios with 100 assets and a 20 day forecasting period and determined that for both weighting methods the RMT filter outperformed the other methods, while the largest eigenvalue filter was the worst performing of the three. Overall, the exponentially weighted, RMT filtered forecasts produced the lowest risk value, by a small margin. Results for other (longer) forecasts were reported to be very similar.

When greater numbers of assets were considered, similar results were found. When the number of assets was reduced, to $N = 50$, historical unfiltered were found to be competitive with RMT filtered forecasts. Finally, Pafka et al. (2004) determined that the unfiltered Riskmetrics (1996) forecasts, using the Riskmetrics (1996) *recommended decay factor*, produced portfolios that were *far riskier* than those produced with filtered forecasts, especially when large numbers of stocks were involved.

The paper of Pafka et al. (2004) was the closest found in the literature to our own questions for a number of reasons. First, they have considered a forecasting period comparable to our own, while also dealing with longer forecasting. Secondly, they have evaluated exponentially weighted, Riskmetrics (1996) style forecasts, which we also wished to consider, particularly in their RMT filtered form. Thirdly, they have averaged over many different test dates, which we considered essential to properly assess the forecasts. Finally, the effect of filtering on portfolios containing different numbers of assets was evaluated. We built on this work in what follows, in particular by assessing the stability-based filter, and including forward validation testing.

4.3 Data

In this chapter we examine the dataset of Daly et al. (2008). These data, used to test filter performance on a stock portfolio, were daily closing prices for the S&P 500 index stocks. The index composition was taken as of 1st February 2006¹. This dataset ran from 1st June 1995 to 1st February 2006, and any series not covering the entire period were discarded, leaving a total of 432 stocks available for testing.

¹from www.standardandpoors.com

4.4 In-sample Testing

4.4.1 Methodology

For the in-sample analysis, and following Pafka et al. (2004), we applied the bootstrapping procedure described in Section 3.3.1. For the S&P 500, we examined portfolios with different numbers of assets (N), considering the cases $N = 100, 250$ and 432 . We also evaluated two forecasting periods, $F = 20$ and 50 days, with 20 days being the one of primary interest, to coincide with other research, such as Pafka et al. (2004) and Morgan and Reuters (1996). We found the results for $F = 20$ and $F = 50$ to be comparable, and concentrated subsequently only on the $F = 20$ case. The filters considered were the LCPB method (Section 3.2.5.1), PG+ method (Section 3.2.5.2), and the KR method (Section 3.2.6.2), with the KR method being our novel filter. We did not consider the filter which retains only the largest eigenvalue, due to its previously measured performance (Pafka et al., 2004).

4.4.2 Measuring Noise

We started by assessing the noise in the unfiltered forecasts, as measured using RMT analysis. Figures 4.1 and 4.2 show, for equally and exponentially weighted forecasts, and for 100 and 432 assets, the percentage of actual measured eigenvalues that were larger than the theoretical maximum eigenvalue predicted by RMT. Both covariance and correlation forecasts were considered. Here, we see that the number of non-noisy eigenvalues increased with both the number of past moves in the equally weighted case, and with the decay factor in the exponential case. In other words, measured in this way, the amount of noise in the unfiltered forecasts was reduced when more data were used. In the case of exponential weights, using a higher decay factor is equivalent to using more data, and here again we saw that this reduced noise.

These effects were expected, following from the work of Pafka and Kondor (2002a,b),

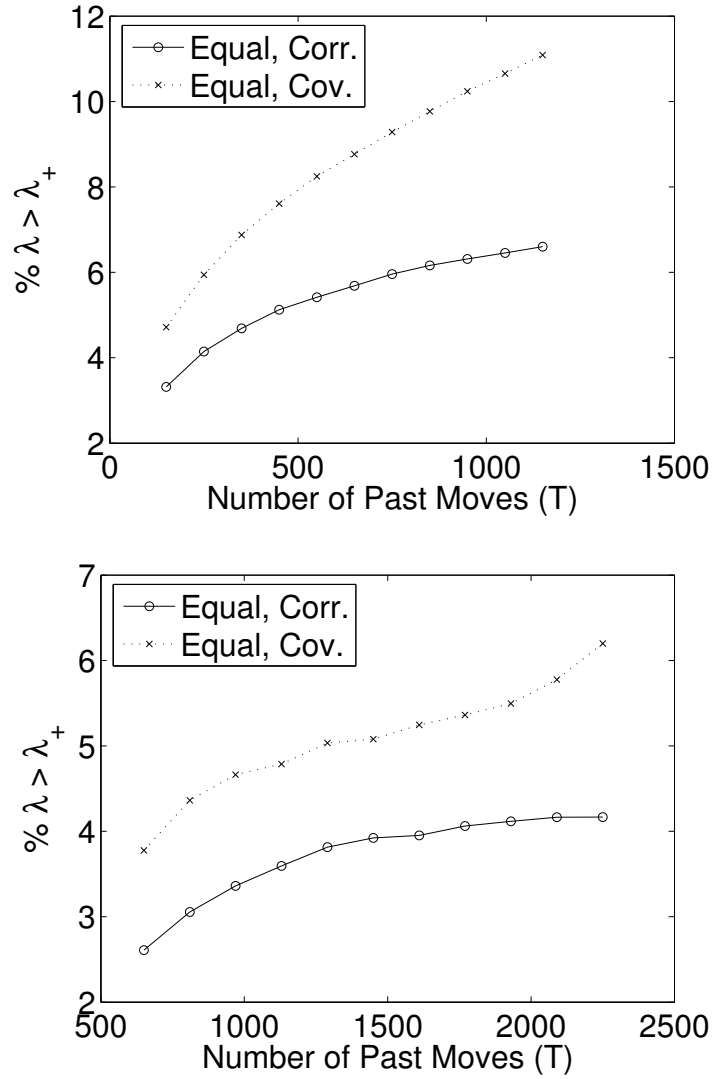


Figure 4.1: Percentage of eigenvalues that were larger than the maximum eigenvalue predicted by RMT, for equally weighted correlation and covariance with 100 assets (above), and 432 assets (below).

which showed that noise was reduced as the ratio of data to assets was increased. However, we note that the suggestion to use more data to reduce noise is not satisfactory, since not only may the data not be available, but long price histories may not be suitable for the medium term forecasting periods we are interested in.

It can also be seen here that, in general, compared to the correlation matrix, the covariance matrix contained more “non-random” eigenvalues. In the case of exponentially

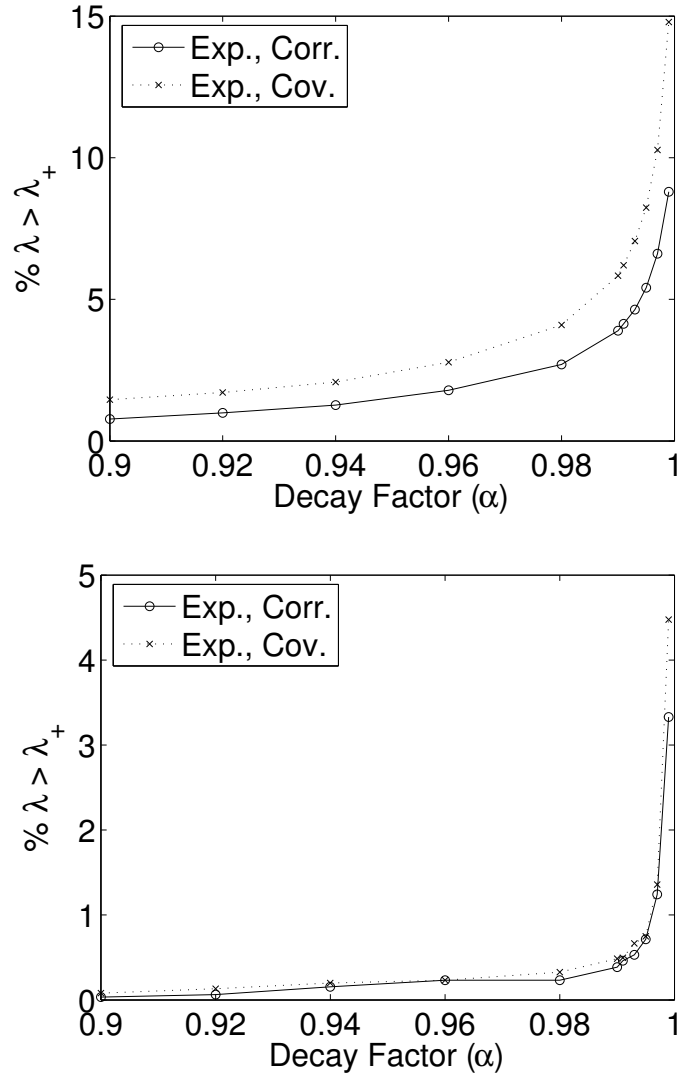


Figure 4.2: Percentage of eigenvalues that were larger than the maximum eigenvalue predicted by RMT, for exponentially weighted correlation and covariance with 100 assets (above) and 432 assets (below).

weighted matrices with 432 assets however, the effect was less pronounced.

Notably, for a wide range of decay factor values, exponentially weighted matrices had very few non-noisy eigenvalues when all 432 assets were used. The range for α included the values of 0.94 and 0.97 suggested by Riskmetrics (1996) for daily and monthly forecasts. Despite this, *filtered* exponentially weighted forecasts produced the lowest mean realised risk in out-of-sample tests, as seen in Section 4.5.

4.4.3 Measuring Stability

For the case of 100 assets, Figures 4.3 and 4.4 display, for selected filters, and averaged over all bootstrap samples, the mean stability across all eigenvectors of the filtered and unfiltered covariance matrices. For these particular graphs equal weights were used, and they show the results for filtering correlation and covariance respectively.

The covariance matrices produced by KR filtering were seen to have better stability than both the LCPB and PG+ filtered ones. This was particularly true for direct filtering of the covariance matrix. It was also seen here that stability improved as the minimum replacement eigenvalue for the KR filter approached zero. Conversely, the closer the minimum eigenvalue approached to the mean noisy eigenvalue, the more stability decreased, although it always remained above that of the LCPB and PG+ filters, as expected. These results were consistent with the definition of the KR filter, which was designed to give improved stability over the other filters.

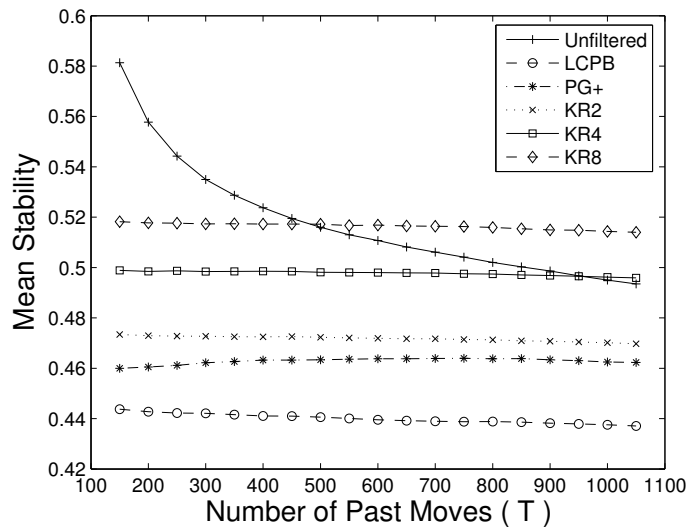


Figure 4.3: In-sample mean stability for the equally weighted covariance forecast with 100 assets, filtering correlation.

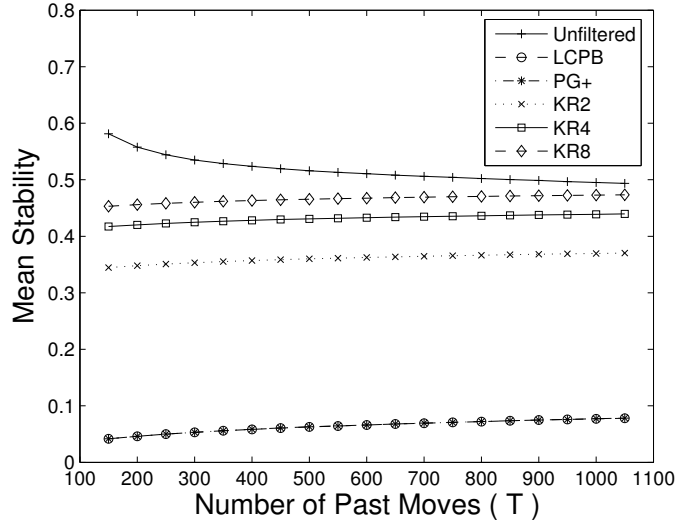


Figure 4.4: In-sample mean stability for the equally weighted covariance forecast with 100 assets, filtering covariance.

4.4.4 In-sample Risk Reduction

Here we review the in-sample results for filtering the stock correlation matrix. The in-sample results showed, in general, the potential of RMT filters to reduce realised risk. Figures 4.5 to 4.9 show the bootstrap results for 20 day forecasting, for equally and exponentially weighted forecasts, for selected filters. We noted some key areas of agreement between our results and those of Pafka et al. (2004), as follows. Firstly, similar risk reduction was achieved, while it is interesting to note that the optimal parameter values suggested were identical in this case. We were also in agreement with Pafka et al. (2004) that longer forecasting periods produced comparable results to the 20 day case, while the results for different numbers of assets were also similar, (with historical unfiltered forecasts being more competitive for the lower numbers of assets). Moreover, RMT filtering performed best for both uniform and exponential weighting, again in agreement with Pafka et al. (2004), and we also found that the Riskmetrics-recommended decay factors were inappropriate for use with the unfiltered forecasts, particularly for the larger numbers of stocks, since they greatly increased realised risk.




Figure 4.5: Mean bootstrapped (in-sample) realised risk, for selected filters, applied to equally weighted volatility forecasts of S&P 500 stocks, and for unfiltered volatility (“ORIG”), for 100 assets. The x-axis shows the Q -value, which is the number of past moves used, divided by the number of assets in the portfolio. Note that filtering is most effective for low values of Q , and produces the best overall risk.

Focusing on the stability-based KR filter, we found that optimisation performance disimproved as the minimum replacement eigenvalue approached zero. The KR method with a minimum replacement eigenvalue of 10^{-8} was not competitive when compared to other methods for reducing risk, (or with the unfiltered series), *despite being the filter with the greatest stability*. In contrast, the LCPB method had the lowest stability, but had reasonably good risk reduction, although not the best. We found a marked risk reduction was achieved *by varying the minimum replacement eigenvalue* of the KR method. We noted, in particular, that the KR2, KR4 and KR8 methods were among the best performing of all filters for this, and were also reasonably consistent with each other.

When considering parameters, in many filtered cases two local minima were produced for the choice of optimal decay factor for the exponential weights ². One coincided with the suggestion of Riskmetrics (1996), i.e. 0.97 for monthly forecasts. The other was much closer to 1. In the case of equally weighted forecasts, with 250 and 432 assets, we noted

²optimal for reducing in-sample realised risk

that the optimal number of past moves was the minimum possible. This was true for many different filters, not just the best one. In contrast, the unfiltered forecasts, when used with this low number of past moves, produced the worst risks seen in the test. This is a very

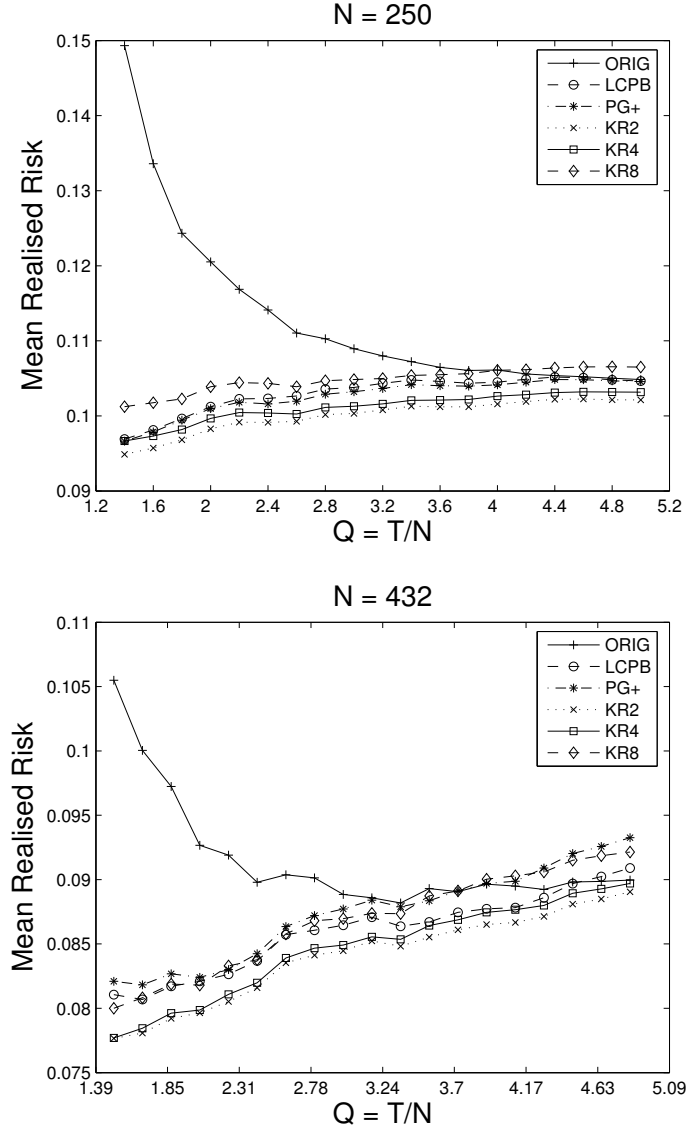


Figure 4.6: Mean bootstrapped (in-sample) realised risk, for selected filters, applied to equally weighted volatility forecasts of S&P 500 stocks, and for unfiltered volatility (“ORIG”), for 250 assets (above) and 432 assets (below). The x-axes show the Q -value, which is the number of past moves used, divided by the number of assets in the portfolio. Filters were more effective, overall, for the larger number of assets, and the best risk occurred, with stability filtering, at the lowest possible value of Q in both cases. Note also the consistency between the stability filters shown.

interesting effect, namely that the best forecasts were found to be ones that reacted most swiftly to market events. Meanwhile, these adaptable methods were not available without filtering.

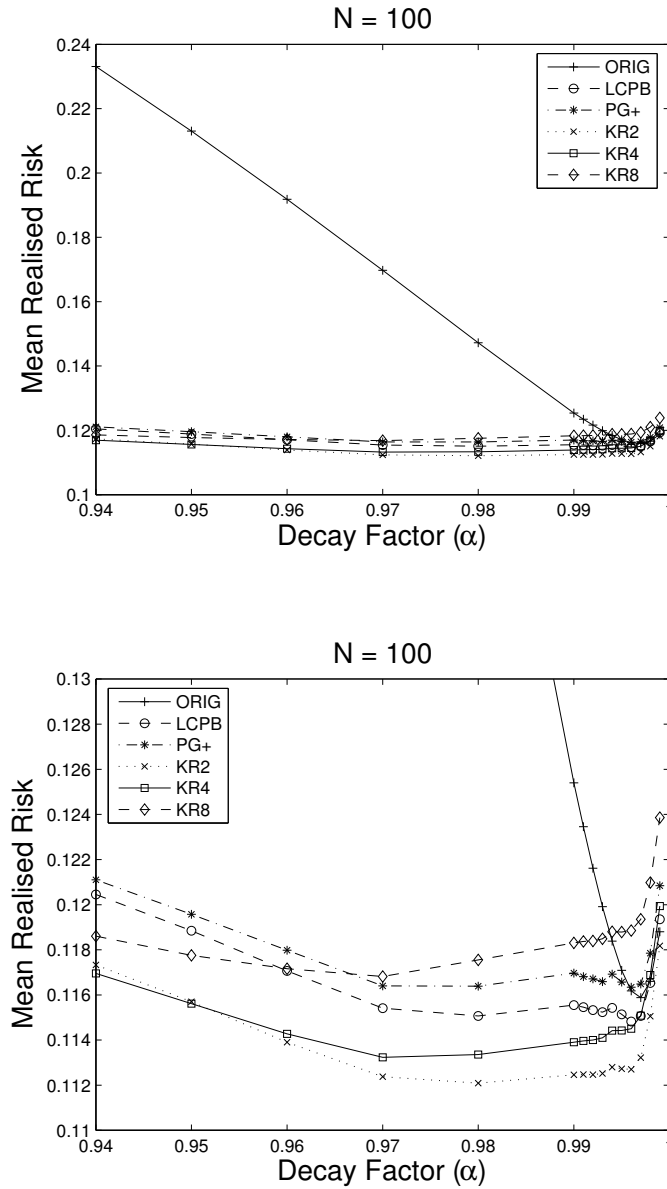


Figure 4.7: Mean bootstrapped (in-sample) realised risk, for selected filters, applied to exponentially weighted volatility forecasts of S&P 500 stocks, and for unfiltered volatility (“ORIG”), for 100 assets (above), and with detailed view (below). RMT filtering is seen to improve risk across a wide range of decay factors, with a stability filter (KR2) resulting in the best overall risk.

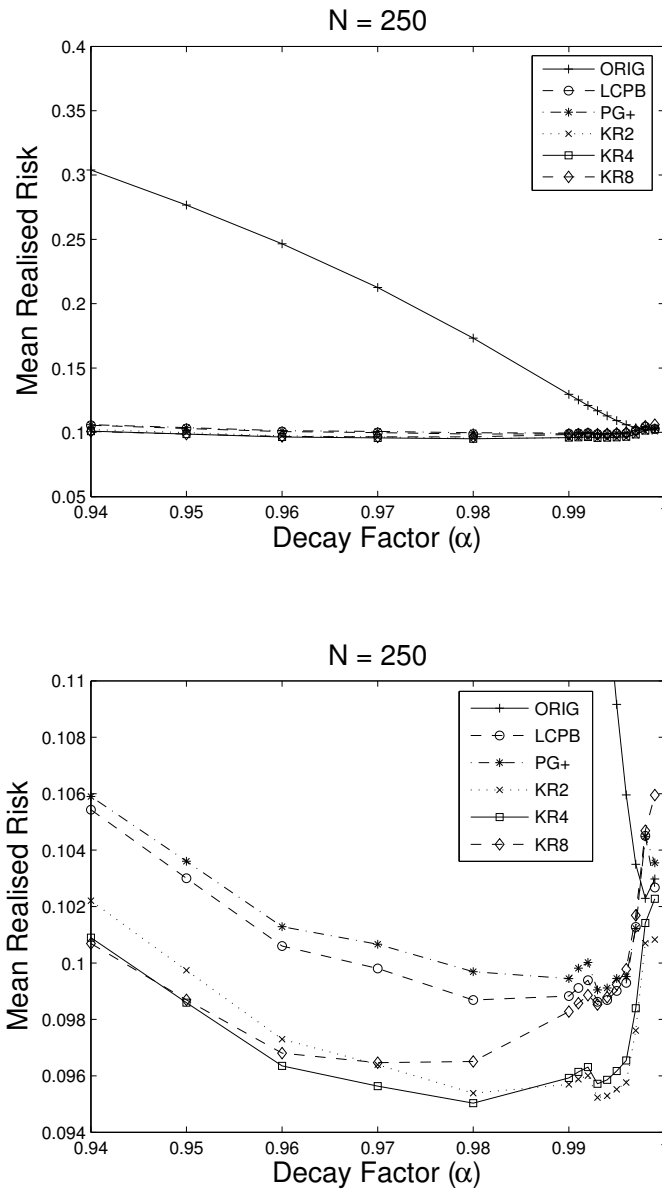


Figure 4.8: Mean bootstrapped (in-sample) realised risk, for selected filters, applied to exponentially weighted volatility forecasts of S&P 500 stocks, and for unfiltered volatility (“ORIG”), for 250 assets (above), and with detailed view (below). The lowest risk was again produced by using a stability filter.

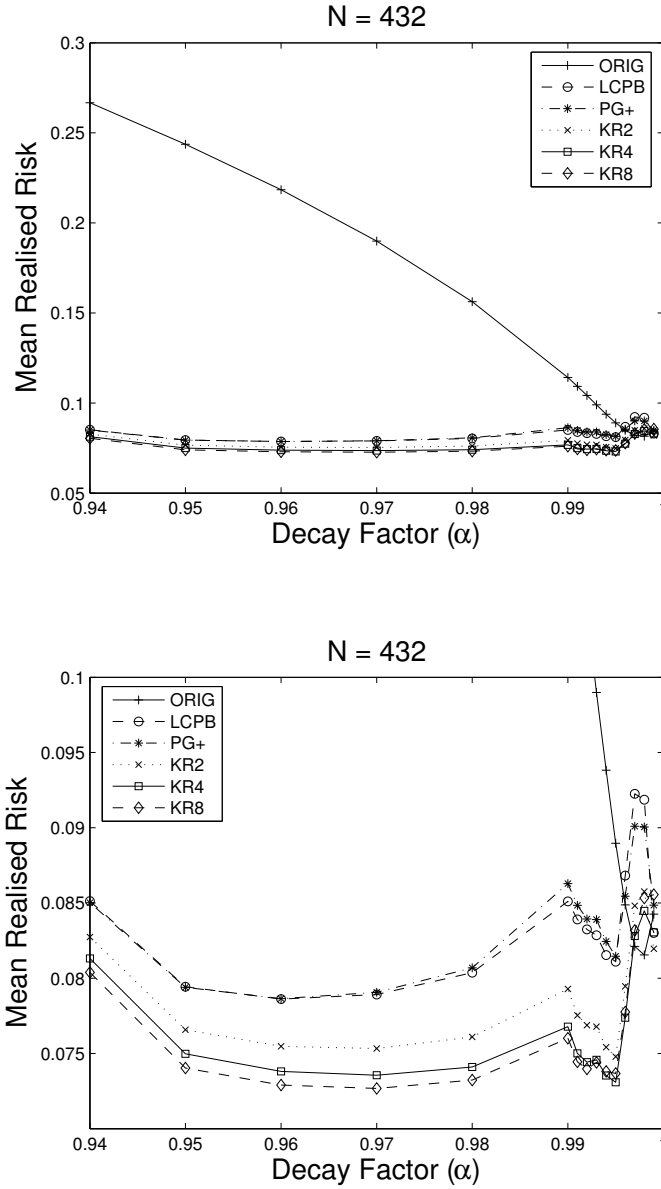


Figure 4.9: Mean bootstrapped (in-sample) realised risk, for selected filters, applied to exponentially weighted volatility forecasts of S&P 500 stocks, and for unfiltered volatility (“ORIG”), for 432 assets (above), and with detailed view (below). Stability filters can again be seen to outperform other models, while two local minima for the decay factor are also seen. Note that the Riskmetrics decay factor of 0.97 performed poorly without filtering (for all three values of N).

4.5 Out-of-sample Testing

4.5.1 Methodology

For comparing the models out-of-sample we used forward validation, as described in Section 3.3.2. To summarise, we considered every available test date and for each one used data prior to the test date to optimise any model parameters. The value of the weighting parameter (α or T) and the choice of KR model were determined out-of-sample. In this case we had 1658 available test days, 129 of which were used as the initial training period. Subsequent retraining was done daily. All 432 assets were used to eliminate the need to arbitrarily choose assets each day.

4.5.2 Out-of-sample Analysis

4.5.2.1 Overall Risk

Table 4.1 shows a summary of the out-of-sample performance of the covariance forecasting and filtering combinations, in the case where correlation was filtered. The figures shown are mean realised risk as a percentage of the result for unfiltered equally weighted covariance. We see that RMT filtering reduced risk *on average* in all cases where it was used. The range of reduction was 12.4% to 14.9% for equal weights, and 5.7% to 10.1% for exponential (these exponential figures are expressed as a percentage of the corresponding unfiltered exponential result. This normalisation occurs in a number of parts of the thesis, (where it has been noted), to facilitate direct comparison with the equally weighted improvements.). The KR filter supplied the most risk reduction in both cases. Overall, the best performing forecast in this test used exponential weights with a KR filter, which had a risk which was 84% of the benchmark.

4.5.2.2 Annual Risk

Table 4.2 shows a breakdown of the results on an “annual” basis over 6 years . In this case, the figures in the table are the mean realised risks, as a percentage of the equally weighted

Table 4.1: Mean out-of-sample realised risk as a percentage of that for unfiltered equally weighted covariance, where correlation was filtered. Filtering was seen to reduce mean realised risk in all cases, and to as low as 84% of our benchmark risk, when applying KR filtering to exponentially weighted forecasts.

Model	Unfiltered	LCPB	PG+	KR
Equal Weights	100	87.3	87.6	85.1
Exponential Weights	93.4	87.1	88.1	84

Table 4.2: Mean out-of-sample realised risk per year for 6 years as a percentage of the equally weighted unfiltered result each year, filtering correlation. In some cases, RMT filtering was seen to **increase** mean realised risk over the course of a year. The KR filter was the only one to reduce realised risk in all cases each year.

Weights	Filter	1	2	3	4	5	6
Equal	None	100	100	100	100	100	100
Equal	LCPB	86.3	86.5	89.8	85.3	94.6	81.3
Equal	PG+	87.1	84.7	89.8	85.3	95.7	83.5
Equal	KR	84.2	84.7	87.4	83.2	91.4	80.2
Exp	None	96.4	96.4	92.9	87.4	95.7	90.1
Exp	LCPB	81.3	82.9	89.8	85.3	94.6	91.2
Exp	PG+	84.9	85.6	89	86.3	91.4	93.4
Exp	KR	80.6	81.1	89.8	80	88.2	85.7

unfiltered result in each year.

Here we see some examples of filtering *increasing* the mean risk in a year. However, the majority of the time filtering reduced risk and, when risk *was* increased, no large increases were observed, the largest being from 90.1 to 93.4, a percentage increase of only 3.7%. Overall, the range of percentage changes was [-19.8%, -4.3%] for equal weights and [-16.4%, +3.7%] for exponential (exponential figures are expressed as a percentage of the corresponding unfiltered exponential result). *The stability-based KR filter was the only filter to reduce mean realised risk in each year in all cases.* For this filter, the range of percentage changes to realised risk was [-19.8%, -8.6%] for equal weights and [-16.4%, -3.3%] for exponential (exponential figures are expressed as a percentage of the corresponding unfiltered exponential result). The overall best method (exponential weighting, KR filter) was found

Table 4.3: Daily frequency of percentage effect, on realised risk, of applying RMT filters to the correlation matrix. “Low” and “High” specify a range for the realised risk, expressed as a percentage of the relevant unfiltered realised risk. Filtering was seen to reduce risk 74.3% of the time, while stability-based filtering reduced it the most often, 80.7% of the time. It can also be seen that, despite reducing mean realised risk overall, all filters had the potential to markedly increase realised risk on any individual day.

Low	High	LCPB	PG+	KR
40	60	38	45	25
60	80	723	649	816
80	100	1431	1464	1626
100	120	668	706	495
120	140	158	153	80
140	160	30	32	11
160	180	7	8	4
180	200	3	1	1

to produce the lowest risk in four of the six years tested, while being competitive in the other two.

4.5.2.3 Daily Risk

Table 4.3 shows the frequency of daily filtering effects. “Low” and “High” specify a range for the realised risk, expressed as a percentage of the relevant unfiltered realised risk. The effects on equally weighted and exponentially weighted matrices were combined here. For example, the LCPB method reduced realised risk to 60%-80% of the unfiltered realised risk for 723 out of 3058 tests, 1529 equally weighted and 1529 exponentially weighted. From this we observe that, taking the mean across all filters, RMT filtering reduced realised risk on 74.3% of the days. The KR method was the most consistent in terms of reducing realised risk, 80.7% of the time overall, compared to 71.7% and 70.6% for the LCPB and PG+ methods respectively.

While filtering reduced realised risk on average, and on the majority of years and days, all the tested filters had the capacity to increase realised risk on any particular day, with some marked increases being observed. Combining all methods, RMT filtering caused an increase in realised risk of 20% or more on 5.3% of the days, while KR filtering increased

realised risk by 20% or more the least often, 3.1% of the days, less than half the frequency of either the LCPB or PG+ methods. On 0.3% of the days (approximately two days every three years), RMT filtering increased realised risk by 60% or more, with the largest increases almost doubling risk. To our knowledge, the capacity of RMT filtering to increase risk in this way had not been previously discussed in the literature. Despite these daily fluctuations, which can be attributed to the arrival of new market information, we note that forecasts using filtering improved risk with good annual consistency.

4.5.3 Parameter Values

Figure 4.10 shows, in the equally weighted case, the number of past moves chosen through time by our forward validation test, for each of the three filtering methods applied to the correlation matrix. This graph shows good agreement with the in-sample tests. For the majority of the forward validation test days, very low numbers of past moves were preferred, which corresponded with the models that were most reactive to market events. The value chosen by the unfiltered series was always the maximum available in this test ($T = 1010$). The reason that a maximum value of T was chosen is so that a consistent “sliding window” of equal length can be used across all tests. The behaviour at different values of T is of interest. The actual maximum value, $T = 1010$, was picked after analysis, to balance the conflicting requirements of (1) capturing as much interesting behaviour as possible, by using a large range for T and (2) performing a forward validation test over the longest possible time frame (the choice of T affects this since the first T days are required for covariance forecasting and cannot therefore explicitly appear in the forward validation). Typically, a minimum of a five year test is considered a good rule of thumb in practice.

Figure 4.11 shows the decay factor values chosen by the forward validation, for each of the three filtering methods applied to the correlation matrix. The decay factor chosen for the unfiltered series, not shown, was always the maximum tested ($\alpha = 0.999$), a value much higher than the Riskmetrics (1996) recommended 0.97. A value of 0.999 yields a markedly different style of model to 0.97, with the former behaving almost like a linear model, as

seen in Figure 4.12. For the filtered forecasts, all the decay factors chosen, using forward validation, were higher than 0.97. We also see here that the optimal parameter values varied over time, as expected.




Figure 4.10: Optimal forward validation number of past moves for equally weighted covariance, filtering the correlation matrix. Note that low values of T were usually chosen for all three filters.




Figure 4.11: Optimal forward validation decay factor values for filtering the correlation matrix. The range of decay factors seen here was higher in all cases than the Riskmetrics (1996) value of 0.97.

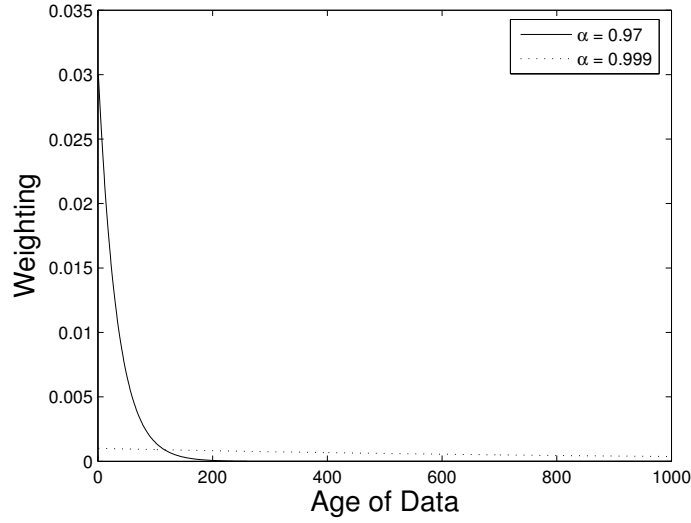


Figure 4.12: First thousand normalised exponential weights corresponding to decay factors of 0.97 and 0.999. Note that the style of models generated are markedly different, with the weights corresponding to 0.999 behaving almost like an equally weighted model.

4.5.4 Range of Realised Risks

In Section 4.5.2 it was observed that filtering had the capacity to *increase* realised risk on individual days. Risk was increased on 25.7% of the days during the forward validation test. In some cases the percentage increases were found to be substantial, with the largest observed in the range of 80% to 100%. Here, we study the effect of these increases on the *range* of the filtered risks, finding that filter performance was not consistent across different risk levels, performing better when unfiltered risk was high. This resulted in the range of the filtered risks being equivalent to, or actually improved upon, that of the unfiltered risk.

Analysing these out-of-sample results further, Figure 4.13 shows the effect of filtering, on each individual day tested, for equally and exponentially weighted covariance respectively. The unfiltered risk is included for comparison, as a straight line. We noted that the bulk of observations were at low unfiltered risk levels. We also observed good risk reduction when the unfiltered series was at the upper end of its range. This last point was particularly true for equal weights, where *all* the filtered risks were below the unfiltered, in the upper

part of the range. At the same time, we observed large percentage increases occurring in the middle and lower ends of the ranges.

Overall, we found that the range of the filtered risks was equivalent to that of the unfiltered risk. For equal weights, the upper limit of the filtered range was actually below that for the unfiltered, while for exponential weights, this was true for two of the three filters (LCPB and KR). For the PG+ filter and exponential weights, the increase to the upper limit of the range was marginal (0.02%).

These results are summarised in Tables 4.4 and 4.5, for equally and exponentially weighted covariance forecasts respectively. Here, we broke the full range of unfiltered risks into quintiles. “All data” refers to the full range, while “Top Quintile” refers to the top 20% unfiltered risks, and so on. The results shown are the average risk, after filtering, in each quintile, as a percentage of the corresponding unfiltered result. Thus, it can be seen that filtering resulted in the *biggest percentage risk reductions during periods where that risk was high*. For equal weights we saw reductions of up to 20%, and for exponential 15.7%, in the top quintile. For all three filters, a general deterioration in performance with each subsequent lower quintile was observed. For the lowest quintiles we saw that some risks were increased after filtering, although not for the KR filter.

To summarise, despite resulting in large increases in individual daily risks, RMT filtering was not found to have an adverse effect on the *range* of realised risk, and so would be a viable tool for use in practical situations.

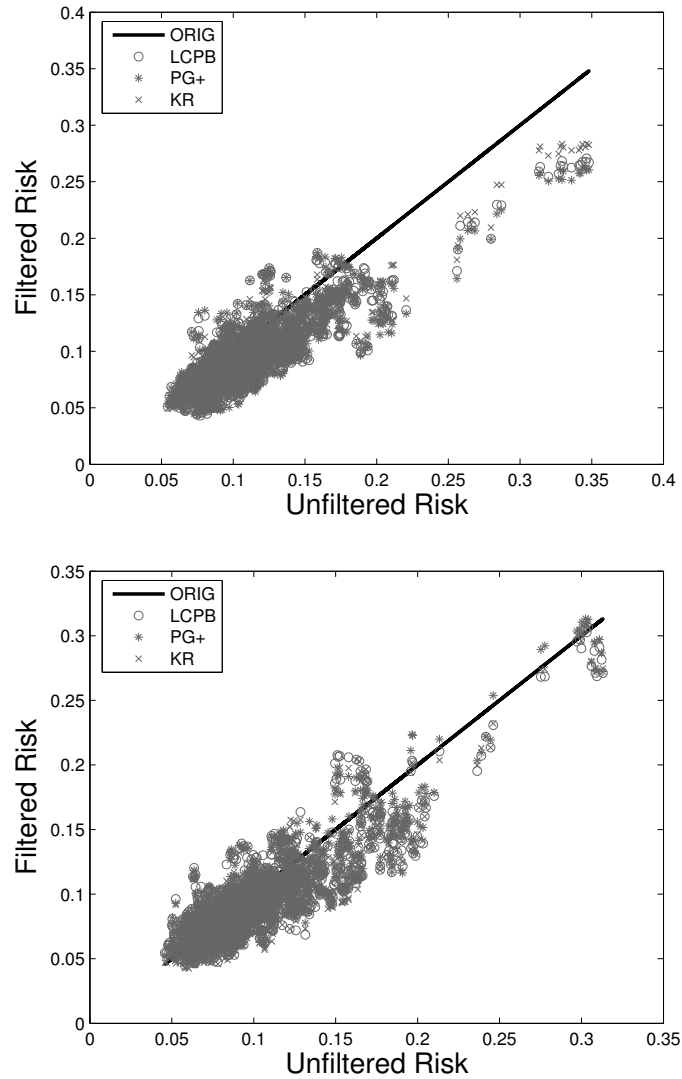


Figure 4.13: Filtered vs. unfiltered realised risks, for the S&P 500 portfolio, for equally (above) and exponentially (below) weighted forecasts. The unfiltered risk is also included, as a 45 degree line, for comparison. It can be seen that filtering was most effective when the unfiltered series was in the upper end of its range, while the large risk increases occurred in the middle and lower ends of this range. Overall the upper limit of the *range* of the filtered risks was either equivalent to, or improved upon, that for the unfiltered risk.

Table 4.4: Summary of filter performance in different unfiltered risk environments, for the S&P 500 portfolio, for equally weighted covariance forecasts. “All data” refers to the full range of unfiltered realised risk. The remaining columns break this range into quintiles. The results shown are the average realised risk in each quintile, as a percentage of the corresponding unfiltered result. Thus it can be seen that the filters were most effective at reducing risk during periods where that risk was high. This resulted in the range of risk being improved upon the unfiltered case for all filters.

Filter	All data	Top Quintile	Quintile 2	Quintile 3	Quintile 4	Bottom Quintile
None	100	100	100	100	100	100
LCPB	87.3	80.3	88.6	90.1	90	94.6
PG+	87.6	80.1	88.3	90.7	90.4	96.2
KR	85.1	80	86.2	86.8	87	90.4

Table 4.5: Summary of filter performance in different unfiltered risk environments, for the S&P 500 portfolio, for exponentially weighted covariance forecasts. This table corresponds to Table 4.4. It can again be seen that the filters were most effective at reducing risk during periods where that risk was high. In this case, this resulted in the range of risk being equivalent to (PG+) or improved upon (LCPB, KR) the unfiltered range.

Filter	All data	Top Quintile	Quintile 2	Quintile 3	Quintile 4	Bottom Quintile
None	100	100	100	100	100	100
LCPB	93.2	84.3	92.8	96.7	96.7	107
PG+	94.3	88.1	93.4	95.4	96.1	107.6
KR	89.9	84.3	89.4	92.6	92	98.6

4.6 Filtering Covariance

In Laloux et al. (2000) and Plerou et al. (2002) the correlation matrix was filtered. While the correlation matrix is closer to the RMT assumption of i.i.d. returns, here we also wished to assess the impact of filtering covariance directly, since it retains more information about the individual assets. Indeed, in Section 4.4.2 we saw that the covariance matrix, when assessed using RMT, was found to have more non-noisy eigenvalues. Here, we determined if this translated into better forecasting performance after filtering. We applied the same methods as used to test filtering correlation, and compared them directly.

4.6.1 In-sample Analysis

Figures 4.14 to 4.18 show the bootstrap results for 20 day forecasting, and equally and exponentially weighted forecasts, for selected filters. The general pattern seen here was very similar to that observed for correlation filtering. RMT filtering still performed best for both uniform and exponential weighting, and the type of improvements seen were comparable with the correlation filter. As was the case with correlation filtering, the results for different numbers of assets were quite similar, with historical unfiltered forecasts being more competitive for the lower numbers of assets. A clear preference for a low number of past moves, in conjunction with filtering, was again observed in the $N = 432$ case. We noted also that the KR2, KR4 and KR8 methods³ were still among the best performing of all filters, and were still consistent with each other.

Some differences to correlation filtering were also noted. For example, different optimal filters were frequently observed in-sample, while the best exponential decay factors were also different. When filtering correlation we observed two competing local minima for the optimal decay factor, one around $[0.97, 0.98]$ and the other in the range $[0.99, 1]$. For direct covariance filtering the latter range was clearly preferred in-sample.

³The KR2, KR4 and KR8 methods are those with minimum replacement eigenvalues equal to the average noisy eigenvalue divided by 2, 4 and 8 respectively.

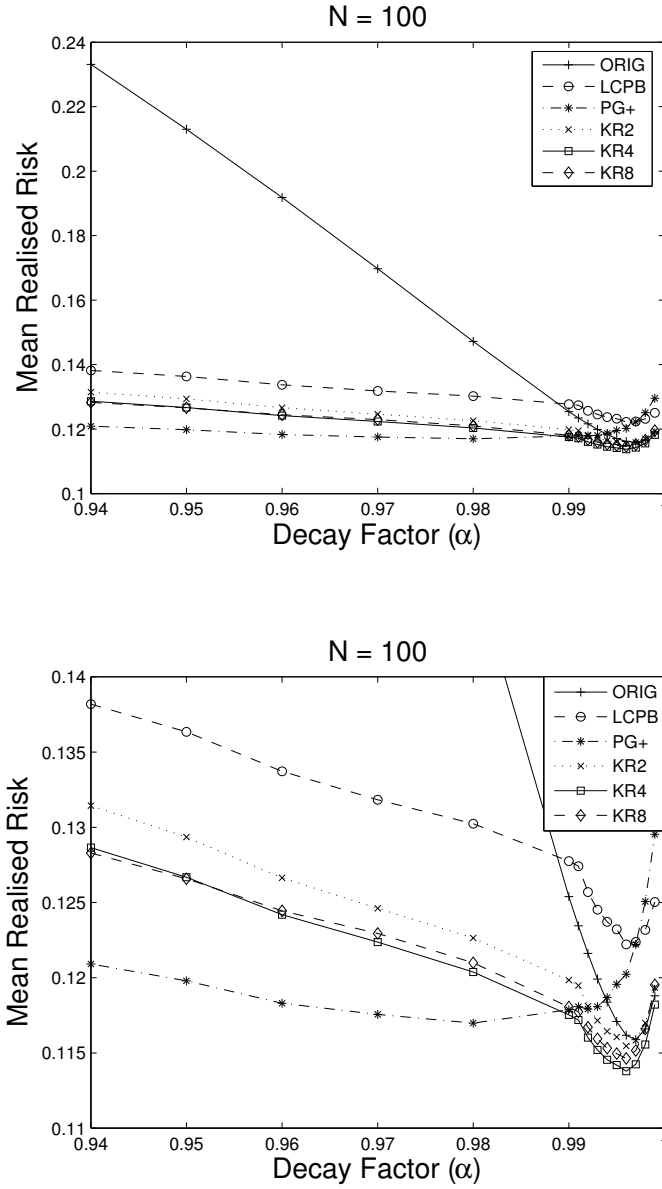


Figure 4.14: Mean bootstrapped (in-sample) realised risk, for selected filters, filtering covariance, with exponentially weighted volatility forecasts of S&P 500 stocks, and for unfiltered volatility (“ORIG”), for 100 assets (above), and with detailed view (below). RMT filtering is again seen to be effective across a wide range of decay factors, with stability filtering being preferred.

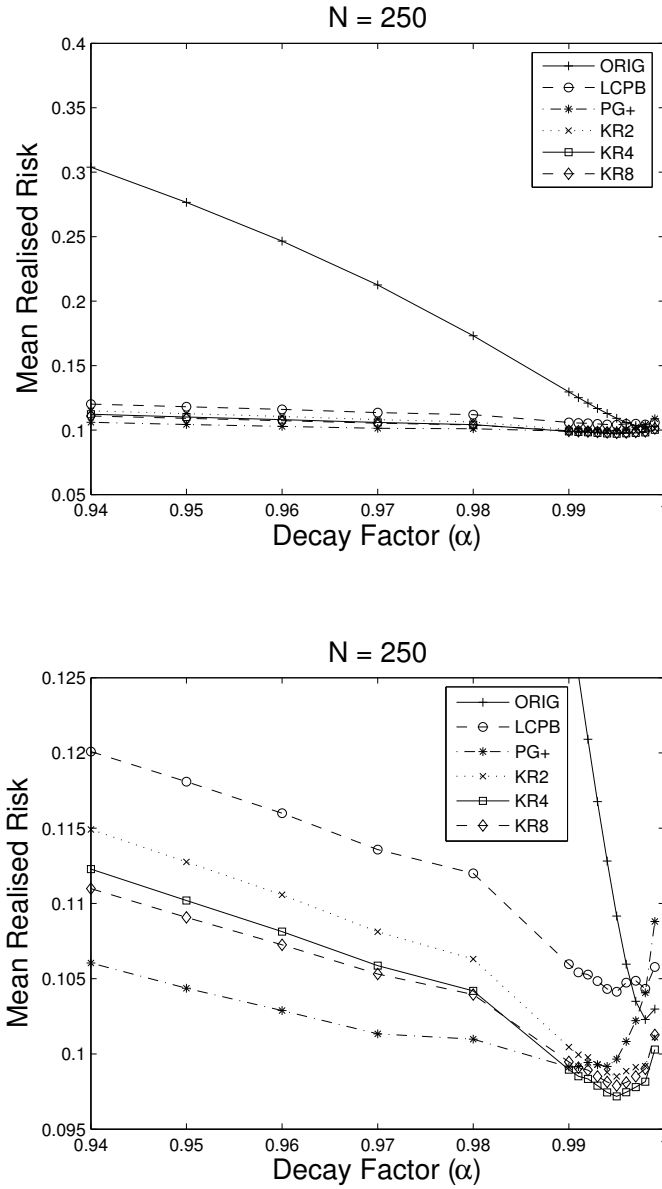


Figure 4.15: Mean bootstrapped (in-sample) realised risk, for selected filters, filtering covariance, with exponentially weighted volatility forecasts of S&P 500 stocks, and for unfiltered volatility (“ORIG”), for 250 assets (above), and with detailed view (below). Stability filtering is again preferred in-sample, with good consistency between the stability filters shown.

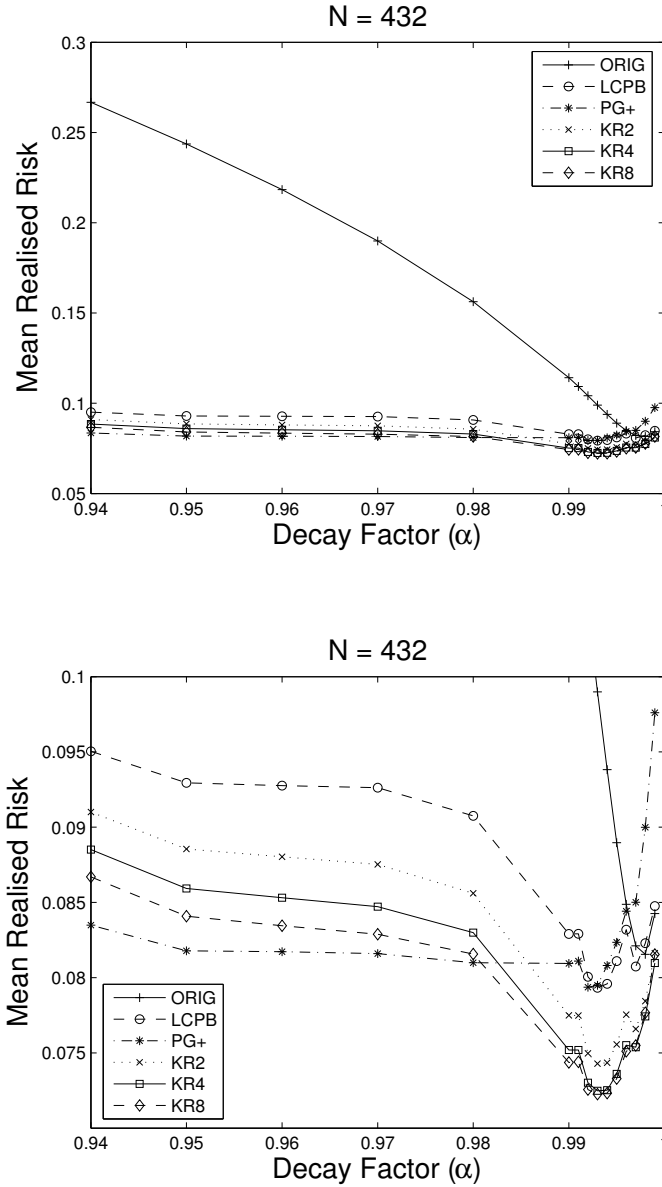


Figure 4.16: Mean bootstrapped (in-sample) realised risk, for selected filters, filtering covariance, with exponentially weighted volatility forecasts of S&P 500 stocks, and for unfiltered volatility (“ORIG”), for 432 assets (above), and with detailed view (below). These figures have much in common with those for filtering correlation, with stability filtering again found to result in the best risks. Here though we see one clear minimum, and subsequent optimal decay factor.

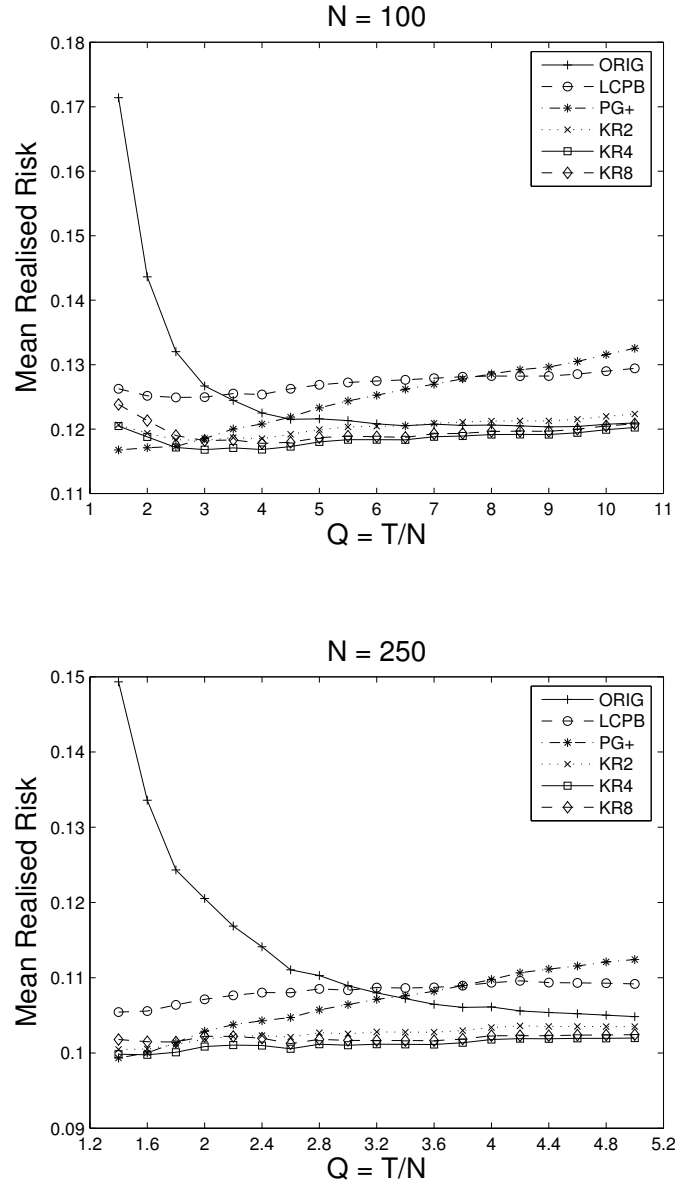


Figure 4.17: Mean bootstrapped (in-sample) realised risk, for selected filters, filtering covariance, with equally weighted volatility forecasts of S&P 500 stocks, and for unfiltered volatility (“ORIG”), for 100 assets (above) and 250 assets (below). The x-axes show the Q -value, which is the number of past moves used, divided by the number of assets in the portfolio.




Figure 4.18: Mean bootstrapped (in-sample) realised risk, for selected filters, filtering covariance, with equally weighted volatility forecasts of S&P 500 stocks, and for unfiltered volatility (“ORIG”), for 432 assets. In this case, an overall preference for the lowest value of Q , with filtering, can again be seen, similar to the correlation case.

4.6.2 Out-of-sample Analysis

We now examine the out-of-sample performance of the direct covariance filtering, for the full portfolio of 432 assets.

4.6.2.1 Overall Risk

Table 4.6 shows a summary of the performance of the forecasting and filtering combinations, in this case. The figures shown are mean realised risk as a percentage of the result for unfiltered equally weighted covariance. RMT filtering was found to reduce risk *on average* in all cases where it was used, as was the case for correlation filtering. Here, the range of reduction was 6.4% to 11.6% for equal weights, and 1.5% to 7.2% for exponential weights (exponential figures are expressed as a percentage of the corresponding unfiltered exponential result). *The KR filter again supplied the most risk reduction for both equal and exponential weights, and the overall best performing forecast in this test was again exponential weighting filtered with the KR method (86.7%).*

Table 4.6: Mean out-of-sample realised risk as a percentage of that for unfiltered equally weighted covariance, with direct filtering of the covariance matrix. Filtering was seen to reduce mean realised risk in all cases, and to as low as 86.7% of the benchmark, for the KR filter applied to exponentially weighted covariance.

Model	Unfiltered	LCPB	PG+	KR
Equal Weights	100	93.6	89.2	88.4
Exponential Weights	93.4	92	87.1	86.7

4.6.2.2 Annual Risk

Table 4.7 shows a breakdown of the mean realised risk of the various weightings and filters on an annual basis. The figures in the table are the mean realised risks, as a percentage of the equally weighted unfiltered result in each year. Similar to correlation filtering, there were a few instances where filtering *increased* the mean risk in a year. However, the majority of the time filtering reduced risk. This annual consistency is an interesting feature, and is seen throughout this work. The range of percentage changes in this case was [-16.8%, 0.0%] for equal weights and [-12.3%, +6.8%] for exponential weights (exponential figures are expressed as a percentage of the corresponding unfiltered exponential result). The stability-based KR filter again reduced mean realised risk in all cases in each year. For this filter, the range of percentage changes to realised risk was [-16.8%, -7.9%] for equal weights and [-12.3%, -2.5%] for exponential (exponential figures are expressed as a percentage of the corresponding unfiltered exponential result), the best ranges of the three filters.

4.6.2.3 Daily Risk

Table 4.8 shows the frequency of daily filtering effects. “Low” and “High” specify a range for the realised risk, expressed as a percentage of the relevant unfiltered realised risk. The effects for equal and exponential weights were combined here. From this we saw that, taking the mean across all filters, RMT filtering reduced realised risk on 70.3% of the days. The KR method was the most consistent in terms of reducing realised risk (80.4% of the time overall, compared to 59.5% and 70.9% for the LCPB and PG+ methods respectively).

Table 4.7: Mean out-of-sample realised risk per year for 6 years as a percentage of equally weighted unfiltered result each year, filtering covariance. In some cases, RMT filtering was seen to **increase** mean realised risk over the course of a year.

Weights	Filter	1	2	3	4	5	6
Equal	None	100	100	100	100	100	100
Equal	LCPB	95.7	100	96.1	87.4	91.4	89
Equal	PG+	90.6	85.6	86.6	87.4	100	87.9
Equal	KR	89.2	91.9	92.1	83.2	87.1	84.6
Exp	None	96.4	96.4	92.9	87.4	95.7	90.1
Exp	LCPB	92.1	95.5	99.2	85.3	84.9	93.4
Exp	PG+	86.3	87.4	91.3	84.2	83.9	90.1
Exp	KR	89.2	89.2	90.6	81.1	83.9	83.5

Table 4.8: Daily frequency of percentage effect, on realised risk, of applying RMT filters directly to the covariance matrix. “Low” and “High” specify a range for the realised risk, expressed as a percentage of the relevant unfiltered realised risk. Filtering was seen to reduce risk 70.3% of the time, while stability-based filtering reduced it the most often, 80.4% of the time. It can again be seen that, despite reducing mean realised risk overall, all filters had the potential to markedly increase realised risk on any individual day.

Low	High	LCPB	PG+	KR
40	60	10	42	5
60	80	449	604	513
80	100	1359	1523	1941
100	120	909	720	575
120	140	278	148	22
140	160	36	15	2
160	180	7	4	0
180	200	8	2	0
200	220	1	0	0
220	240	1	0	0

Combining all methods, RMT filtering caused an increase in realised risk of 20% or more around 5.7% of the days. KR filtering increased realised risk by 20% or more by far the least often (0.8% of the days, compared to 10.8% for LCPB and 5.5% for PG+). On 0.3% of the days, RMT filtering increased realised risk by 60% or more, more than doubling it on two of those days. Thus we again see that consistent daily forecasting was not achieved, and that it was not required to achieve consistent annual improvement.

4.6.3 Parameter Values

Figure 4.19 shows, in the equally weighted case, the number of past moves chosen through time by our forward validation test, for each of the three filtering methods applied to the covariance matrix. Here we saw a different pattern from the correlation filtering case, as the lowest values are not preferred, in particular by the KR method. The value chosen by the unfiltered series was always the maximum available in this test ($T = 1010$).

Figure 4.20 shows the decay factor values chosen through time by the forward validation in this case. The decay factor chosen for the unfiltered series, not shown, was always the maximum tested ($\alpha = 0.999$). All the decay factors chosen using forward validation were higher than the 0.97 suggested by Riskmetrics (1996), and the one for the preferred KR filter reached the highest values of all the filters. Similar to correlation filtering, we saw that the optimal decay factors, like the number of past moves, varied over time. In this case the range of decay factors was higher than the correlation case.




Figure 4.19: Optimal number of past moves for equally weighted covariance, filtering the covariance matrix directly. Here we see that the best parameter choices varied across the range, over time.




Figure 4.20: Optimal forward validation decay factor values for filtering the covariance matrix directly. The range of decay factors seen here was higher in all cases than the Riskmetrics (1996) value of 0.97.

4.6.4 Range of Realised Risks

In this section we review the range of realised risks in the covariance filtering case. As can be seen from Figure 4.21, a similar pattern to the correlation filtering case was observed. For equal weights, the range was reduced after filtering, and the large percentage increases were all found in the lower and middle parts of the range. This information is summarised in Table 4.9 where we see risk reduction by quintile. Again, the percentage effect of filtering improved in general with higher quintiles. In this case, the KR filter was very consistent across quintiles.

For exponential weights, we also saw the largest percentage increases in the lower and middle of the range. However, here we saw an overall increase in the upper limit of this range for the LCPB and PG+ filters, but not for the KR filter. The KR filter was also the only one to reduce in all quintiles for exponential, as seen in Table 4.10, where it was again very consistent. In this table, we see the same general trend of better percentage reduction in higher quintiles.

Overall, RMT filtering cannot be said to have adversely affected the range of realised

risk, despite the large individual daily percentage increases that were observed, since these increases were concentrated in areas where unfiltered risk was low to medium.

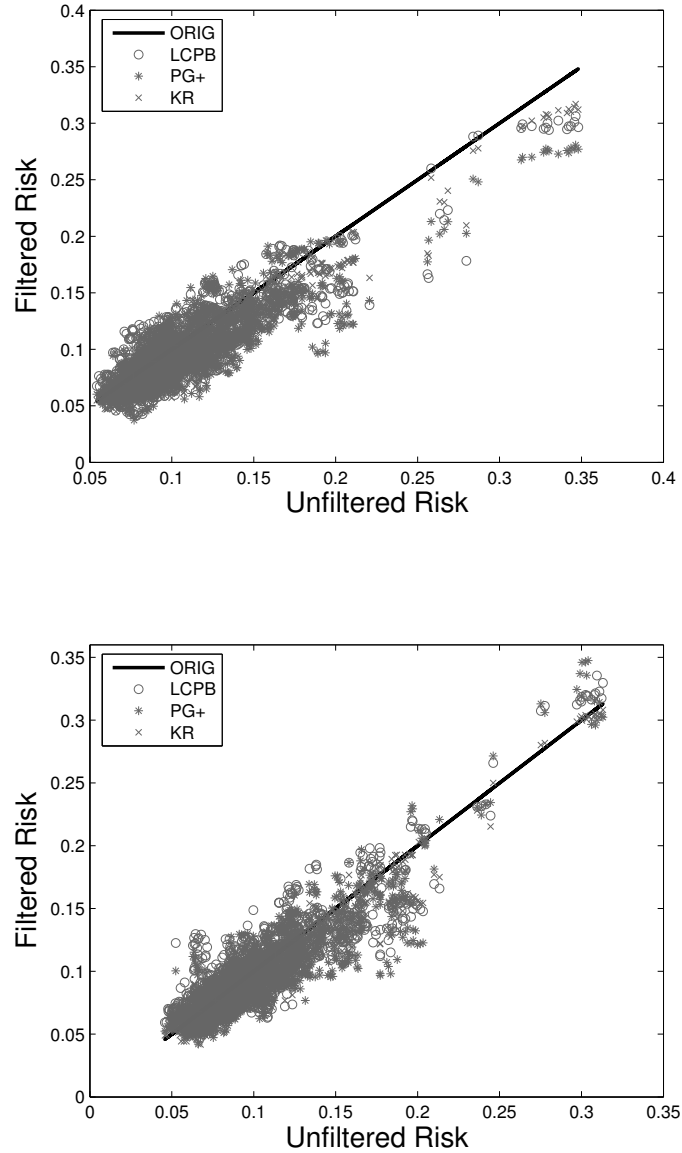


Figure 4.21: Filtered vs. unfiltered realised risks, for the S&P 500 portfolio, for equally (above) and exponentially (below) weighted forecasts, using direct covariance filtering. The unfiltered risk is also included, as a 45 degree line, for comparison. It can be seen that filtering was most effective when the unfiltered series was in the upper end of its range, while the large risk increases occurred in the middle and lower ends of this range. For equal weights, the upper limit of the *range* of the filtered risks was less than that for the unfiltered risk, while, for exponential, some increase in the upper limit was seen.

Table 4.9: Summary of filter performance in different unfiltered risk environments, for the S&P 500 portfolio, for equally weighted covariance forecasts, using direct covariance filtering. “All data” refers to the full range of unfiltered realised risk. The remaining columns break this range into quintiles. The results shown are the average realised risk in each quintile, as a percentage of the corresponding unfiltered result. Thus it can be seen that the filters were most effective at reducing risk during periods where that risk was high.

Filter	All data	Top Quintile	Quintile 2	Quintile 3	Quintile 4	Bottom Quintile
None	100	100	100	100	100	100
LCPB	93.6	88.8	92.2	94.8	96.3	102.3
PG+	89.2	82.5	90.6	92.2	91.5	95.7
KR	88.4	89	88.2	87.9	86.9	89.6

Table 4.10: Summary of filter performance in different unfiltered risk environments, for the S&P 500 portfolio, for exponentially weighted covariance forecasts, using direct covariance filtering. This table corresponds to Table 4.9. We again see that the filters were most effective at reducing risk during periods where that risk was high.

Filter	All data	Top Quintile	Quintile 2	Quintile 3	Quintile 4	Bottom Quintile
None	100	100	100	100	100	100
LCPB	98.5	93.2	94.2	101	101.4	111.9
PG+	93.3	89.4	91.8	95.3	93.8	102.1
KR	92.9	92.4	92.9	92	92.6	95.6

4.7 Comparison of Tested Filters

These results showed that direct covariance filtering was viable, and moreover reduced overall risk in all cases. We now compare the results from covariance and correlation filtering directly, and also consider the best choice of weighting system in this test.

4.7.1 In-sample Comparison

Figure 4.22 shows an extract of the in-sample results, comparing the best, or two best, filtering methods, in each of the following cases: equally weighted filtered (i) correlation and (ii) covariance, and exponentially weighted filtered (iii) correlation and (iv) covariance. Results are shown (on the same graph) vs. decay factors in the case of exponential weighting and vs. number of past moves in the case of equal weighting. The legends can be interpreted as “forecast weighting (equal or exponential), matrix filtered (correlation or covariance), filtering method”. For example, “Equal, Corr, KR2” refers to the mean realised risk over all bootstrapped samples for equally weighted forecasts, filtered using method KR2 on the correlation matrix.

In the case of 100 assets, a correlation filter achieved lowest risk for both equal and exponential weights. In both cases a KR filter was preferred, while there was little to choose between the two types of weighting. Moving to 432 assets, we found direct filtering of the covariance matrix was better in-sample, in both cases. Again, a KR filter was best, while in this case exponential weighting achieved the lowest risk.

4.7.2 Out-of-sample Comparison

We now consider the out-of-sample results. Table 4.11 shows the percentage of times that best performance was achieved by each method in the forward validation test, on an annual, monthly and daily basis. One month was assumed equal to exactly 21 trading days for this purpose. The daily results showed that an unfiltered forecast was best for only 6% of days, and on the majority of these days the best unfiltered forecast was exponentially weighted.


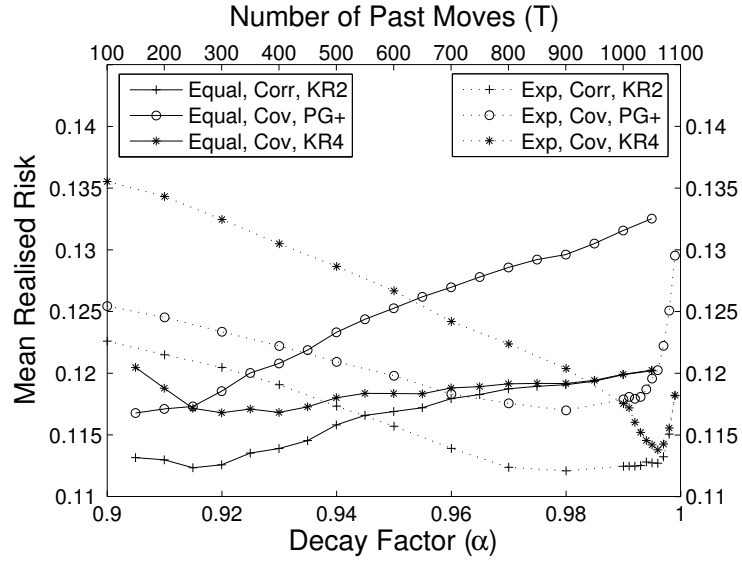


Figure 4.22: In-sample mean realised risk for selected best forecasting methods, for 100 assets (above) and 432 assets (below), showing both equal and exponential weighting schemes. Stability filters are preferred in all cases.

Table 4.11: Percentage of the time each method had the lowest mean out-of-sample realised risk on an annual, monthly and daily basis. The most consistent method by far was the stability-based filter, applied to the correlation matrix of the exponentially weighted forecasts.

Weights	Matrix Filtered	Filter	Yearly	Monthly	Daily
Equal	-	None	0	0	0.8
Exp	-	None	0	5.5	5.2
Equal	Corr	LCPB	0	4.1	7.1
Equal	Corr	PG+	0	9.6	6.9
Equal	Corr	KR	16.7	8.2	7.7
Exp	Corr	LCPB	0	6.2	4.9
Exp	Corr	PG+	0	4.1	6.1
Exp	Corr	KR	50	25.3	19
Equal	Cov	LCPB	0	4.1	3.3
Equal	Cov	PG+	16.7	8.2	7.5
Equal	Cov	KR	0	2.7	4.6
Exp	Cov	LCPB	0	6.8	6.9
Exp	Cov	PG+	16.7	8.2	8.8
Exp	Cov	KR	0	6.8	11.2

The fact that unfiltered forecasting was found to be best for 5.5% of the months reflects some clustering of these daily effects.

The overall best method out-of-sample, KR filtering of the exponentially weighted correlation matrix, was consistently best on an annual (50%), monthly (25.3%) and daily (19%) basis. A covariance filtering method was best 33.3% of the years, 36.8% of the months and 42.3% of the days. This means that filtering the correlation performed best the majority of the time.

Table 4.12 displays a direct comparison between the overall risk of the different methods, sorted from best to worst. Here we can see that the three best methods involved KR filtering, while in the top two it was the correlation matrix that was filtered. For both types of weighting, the KR filter was the best risk reducer, while four of the top five methods involved exponential weights. Overall, our results suggested that correlation filtering was preferred over direct covariance filtering.

Table 4.12: Mean out-of-sample realised risk as a percentage of that for unfiltered equally weighted covariance, for all methods, and sorted best to worst. The three best methods involved KR filtering, while four of the best five used exponential weights.

Weights	Matrix Filtered	Filter	Mean Risk
Exponential	Corr	KR	84
Equal	Corr	KR	85.1
Exponential	Cov	KR	86.7
Exponential	Corr	LCPB	87.1
Exponential	Cov	PG+	87.1
Equal	Corr	LCPB	87.3
Equal	Corr	PG+	87.6
Exponential	Corr	PG+	88.1
Equal	Cov	KR	88.4
Equal	Cov	PG+	89.2
Exponential	Cov	LCPB	92
Exponential	-	None	93.4
Equal	Cov	LCPB	93.6
Equal	-	None	100

4.8 Summary

In this chapter we have examined the application of three RMT filters, including one novel filter, to the optimisation of an S&P 500 portfolio. We have implemented the bootstrapping technique of Pafka et al. (2004) to compare the filters in-sample, and have supplemented this with a full, out-of-sample, forward validation test. Broadly, our results are in agreement with previous results (Pafka et al., 2004), that RMT-based filtering can improve the realised risk of minimum risk portfolios.

In-sample, our bootstrap tests for the standard filters were in close agreement with Pafka et al. (2004). Here, we found similar risk profiles and the choice of optimal parameters was found to be identical. We also determined that the results for our two forecasting periods, 20 and 50 days, were comparable, and found that the competitiveness of unfiltered forecasts improved as the number of assets was decreased. It was also observed that the Riskmetrics-recommended decay factor was unsuitable for use with the unfiltered forecasts.

Subsequently, we determined that, with filtering, two alternative ranges for the optimal decay factor were viable. For equal weights, we observed that, in particular for 250 and 432 assets, forecasts which used a very low number of past moves were optimal, indicating *reactive models*. Without filtering these models were found to be hidden by noise.

While in-sample analysis provides valuable information, about filter behaviour across a wide range of parameters for example, we used out-of-sample forward validation to simulate an actual implementation, requiring daily parameter tuning. Out-of-sample, we studied the effect of filtering on an overall, annual, monthly, and daily basis. RMT filters were found to reduce mean realised risk, overall, in *all* cases tested. However, in some individual years this was not the case. When considering individual days, RMT filtering was found to reduce realised risk for 74.3% of the test cases. However, it was also found capable of increasing realised risk for all types of filters, substantially in some cases. Overall, RMT filtering of the correlation matrix was found to reduce realised risk by between 12.4% and 14.9% for equal weights, and by 5.7% to 10.1% for exponential.

We also examined the behaviour of the RMT filters at different levels of the unfiltered realised risk, with a view to establishing the effect of filtering on the *range* of the realised risk. In general, we found that the days when risk was increased were more likely when unfiltered risk was in the lower and middle parts of its range, while in the lower quintiles average risk was frequently increased. We observed good risk reduction when the unfiltered series was in the upper end of its range. In the case of equal weights, all filtered risks in the upper part of the range were below the unfiltered risk. As a result, the range of the filtered risks was equivalent to the unfiltered range for exponential weights, while for equal weights the upper limit of the range was *reduced* after filtering.

In-sample tests supplied some evidence, in the form of local optima, to support the Riskmetrics (1996) recommended decay factor of 0.97. However, the optimal out-of-sample decay factors, for both filtered and unfiltered forecasts, were higher *in all cases* than those suggested by Riskmetrics (1996), with those for the latter approaching a value of $\alpha = 1$. In practice, such a system, with a decay factor approaching 1, behaves almost like an equally

weighted model.

When considering the covariance matrix, we found that it contained more non-noisy eigenvalues than the correlation matrix. When filtering the covariance directly was examined in-sample, it produced lower risk portfolios than correlation filtering in some cases. Out-of-sample, covariance filtering reduced in all cases, but on average filtering correlation generated lower realised risk. Covariance filtering was also found to result in a lower frequency of daily risk reduction, while the range of increases observed included larger values. Overall, we determined both methods to be viable, but correlation filtering to be preferable, based on these tests.

On comparing our novel filter to the two well-known filters, we found that it offered improvements in terms of risk and stability. When filtering correlation we found that, for equally weighted forecasts, our filter produced a 17.3% improvement in risk reduction compared to the *best* of the other filters. When moving to exponential weights this improvement rose to 49.2%. We found that, as well as contributing the best overall risk for both types of weighting, the KR filter had the highest frequency of risk reduction, at 80.7%. The next highest frequency was 71.7% for the LCPB filter. In addition, the KR filter resulted in only 3.1% of the days having risk increased by more than 20%, less than half the frequency of the other filters.

The KR filter was also the only filter to reduce risk in every year tested, and out of these twelve annual reductions (six for equal weights and six for exponential), only once did another filter outperform the stability-based filter. The overall best forecasting method, out-of-sample, was exponentially weighted covariance, with our Krzanowski stability-based filter applied to the correlation matrix, which resulted in a total risk reduction to 84% of the benchmark. Finally, the KR filter performed consistently well across the full range of unfiltered risks during this test. It returned the lowest risk in all five quintiles for both equal and exponential weights, while in the exponential case it was the only filter to reduce risk in each quintile.

In the next chapter, we study the application of these filters to a foreign exchange and

commodity portfolio with only 39 assets, and report many similar results. This asset class was chosen to utilise previous hedge fund industry experience, trading with systematic mathematical models, especially of foreign exchange.

Chapter 5

RMT for Foreign Exchange Portfolios

5.1 Introduction

In the previous chapter we investigated the behaviour of RMT filters applied to a large stock portfolio, where the amount of data was limited in comparison to the number of available assets. In our case, the lowest number of stocks considered was 100. In this chapter, we extend this analysis to filter correlation matrices of a foreign exchange and commodity portfolio, where the number of tradeable assets was smaller again than the stock market. We considered a portfolio containing *at most* 39 assets. The ratio of available data points to the number of assets was thus higher than the previous chapter, and we wished to assess the benefit of noise filtering in this data rich environment. We noted that, as with stocks, older data may not be relevant for our forecasting period, and thus noise filtering may be successful in reducing portfolio risk.

We first assessed the accuracy of the RMT limiting approximations for such a low number of assets. Finding this satisfactory, we measured the amount of noise in the correlation matrices, using RMT analysis. We then examined the filter behaviour in-sample. We started by reducing the number of assets in the S&P 500 portfolio to 39, for comparison purposes.

Table 5.1: List of currencies and commodities used. The base currency was the U.S. Dollar.

Name	Code	Name	Code
Australian Dollar	AUD	Peruvian Nuevo Sol	PEN
Brazilian Real	BRL	Philippine Peso	PHP
Canadian Dollar	CAD	Platinum Ounce	XPT
Chilean Peso	CLP	Polish Zloty	PLN
Colombian Peso	COP	Romanian New Leu	RON
Czech Koruna	CZK	Russian Rouble	RUB
Euro	EUR	Silver Ounce	XAG
Fijian Dollar	FJD	Singapore Dollar	SGD
Gold Ounce	XAU	Slovak Koruna	SKK
Hungarian Forint	HUF	South African Rand	ZAR
Icelandic Krona	ISK	South Korean Won	KRW
Indian Rupee	INR	Sri Lankan Rupee	LKR
Indonesian Rupiah	IDR	Pound Sterling	GBP
Israeli New Sheqel	ILS	Swedish Krona	SEK
Japanese Yen	JPY	Swiss Franc	CHF
Mexican Peso	MXN	Thai Baht	THB
Moroccan Dirham	MAD	Thaiwanese Dollar	TWD
New Zealand Dollar	NZD	Tunisian Dinar	TND
Norwegian Krone	NOK	Turkish New Lira	TRY
Oil, Brent Crude, Barrel	XCB		

We then examined the 39 asset Fx portfolio and, finally, we reduced the assets in this portfolio even further to assess the limits on filtering. Out-of-sample, we examined the full 39 asset case, as well as a portfolio consisting of 15 major assets from this portfolio. For the full portfolio, we also measured the optimal parameter values, and the effect of filtering on the range of realised risk.

5.2 Data

To examine the application of RMT filters to the currency and commodity portfolio we considered the dataset of Daly et al. (2009). This consisted of daily currency spot prices vs. the US Dollar (USD), along with equivalent rates for Silver, Gold, Platinum and Oil. This data, provided by Pacific Exchange Rate Service¹, covered the period from 4th January 1999 to 31st December 2007. The currencies and commodities selected for this analysis are outlined in Table 5.1. We chose the largest group possible for this, while discarding currencies which were unsuitable due to, for example, regulatory or political restrictions on trading. This group of assets is referred to as the “Fx portfolio” throughout this thesis.

In this chapter we also assess the effect of further reducing the number of assets in the S&P 500 portfolio, studied in Chapter 4, to a size comparable with the Fx portfolio. For this we have reused the dataset of the original S&P 500 analysis, outlined in Section 4.3.

5.3 Accuracy of the Limiting Approximations for 39 Assets

As a preamble, we measured here the accuracy of the RMT limiting approximations, in comparison with simulated random returns for 39 assets. Here we have chosen two typical parameter values, 40 past days for equal weights and a decay factor of 0.97 for exponential. Figure 5.1 shows a sample comparison, between the eigenvalue distribution of 1000 sampled random matrices, and the corresponding RMT approximation, for equally and exponentially weighted matrices.

¹<http://fx.sauder.ubc.ca/data.html>

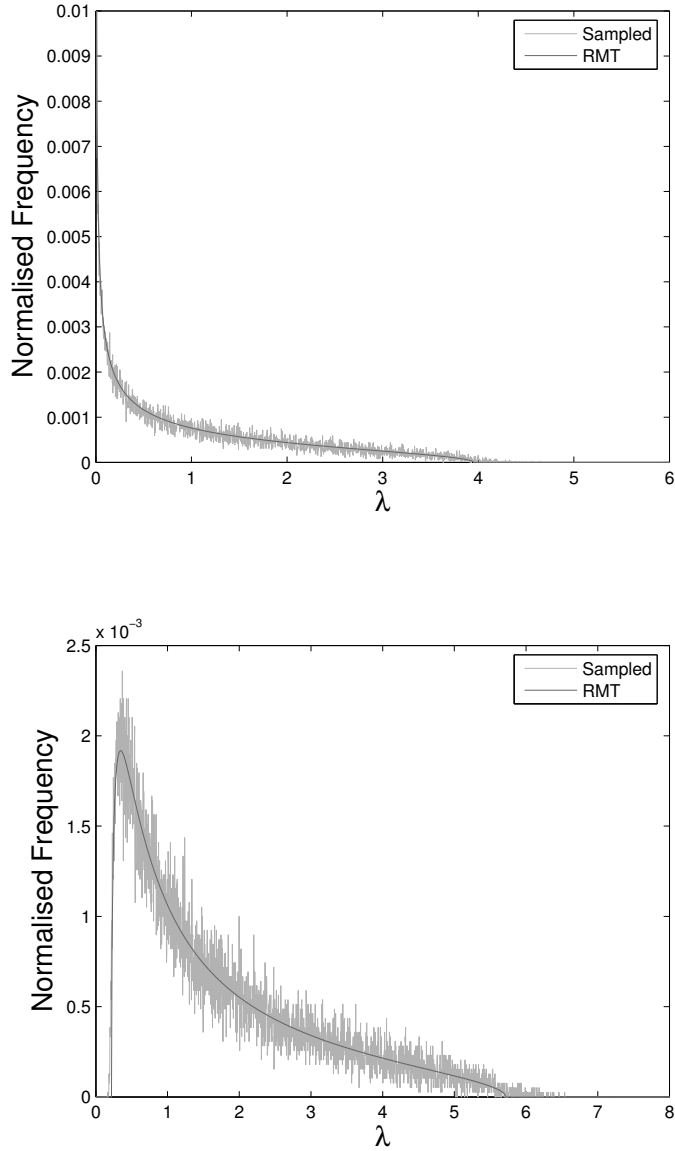


Figure 5.1: Eigenvalue distribution of 1000 sampled random matrices vs RMT approximation, for 39 assets, with equal weights and $T = 40$ historical data points (above), and with exponential weights and a decay factor of $\alpha = 0.97$ (below).

We saw good agreement between the sampled distributions and their RMT approximations, while noting that the RMT distributions, which are approximations as $N \rightarrow \infty$, underestimated the maximum random eigenvalue in both sample cases. This may result in slightly more eigenvalues being identified as containing information, when using the

limiting approximation. Adjustments for finite numbers of assets are discussed further in Chapter 6, under future work.

5.4 Measuring Noise

In this section we assess the noise in the unfiltered forecasts, as measured using RMT analysis. This can be compared to the same analysis, performed for the stock portfolio, in Section 4.4.2. Figure 5.2 shows, for equally and exponentially weighted forecasts, and for 39 assets, the percentage of actual measured eigenvalues that were larger than the corresponding theoretical maximum eigenvalue predicted by RMT. This showed good agreement with the stock portfolio case.

We saw again that the number of non-noisy eigenvalues increased with both the number of past moves in the equally weighted case, and with the decay factor in the exponential case. However, as before, we note that using long data histories to reduce noise can be unsatisfactory, since the data can be unavailable, or simply not relevant to medium term forecasting.

We observed that, as in the S&P 500 case, exponentially weighted matrices had very few non-noisy eigenvalues when all 39 assets were used, for a wide range of decay factors. The range for α included values of 0.94 and 0.97, suggested by Riskmetrics (1996) for daily and monthly forecasts.

5.5 In-sample Testing

5.5.1 Methodology

For the in-sample analysis we repeated the methodology which was applied to the S&P 500 portfolio, outlined in Section 3.3.1. To summarise, for a given value of N (the number of assets), we randomly selected N assets from the data set, and a random test date. For each N , we repeated this random selection 1000 times, with replacement, and calculated the mean, across all of these samples, of the realised risk of the forecast minimum risk

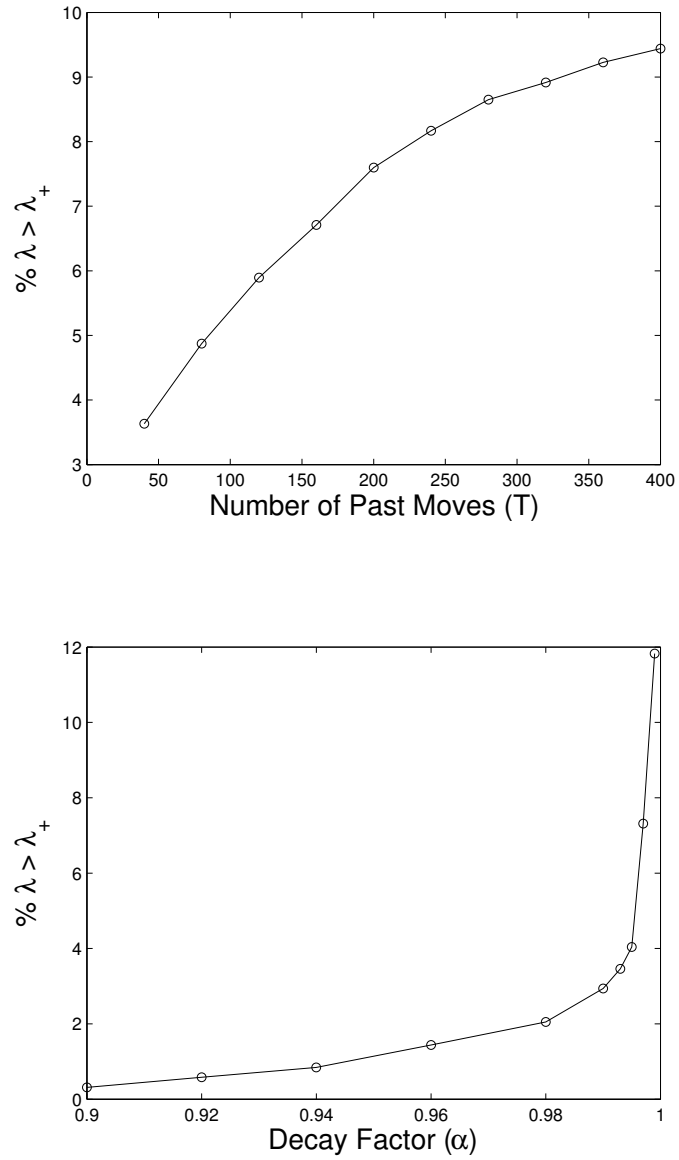


Figure 5.2: Percentage of eigenvalues that were larger than the maximum eigenvalue predicted by RMT, for equally (above) and exponentially (below) weighted covariance, with 39 assets. As the amount of past data is increased, (corresponding to an increase of decay factor in the exponential case), the amount of non-noisy eigenvalues is increased. However, using long price histories may not lead to the best medium term forecasts.

portfolio. This is the bootstrapping procedure of Pafka et al. (2004), applied to our data and filters.

In this section, we first review in-sample results for S&P 500 portfolios with 39 assets, so that we can compare these with the Fx portfolio. In the foreign exchange case, we considered portfolios with the full 39 assets. Subsequently, we also examined Fx portfolios with 10, 15, 20 and 30 assets, to study the limitations of RMT filtering in this case. A covariance forecast again consisted of a raw forecast, which was either exponentially or equally weighted, and could be unfiltered, or filtered by one of the LCPB, PG+ or KR methods, applied to the correlation matrix. In this chapter we have considered a forecasting period of $F = 20$ days for our analysis.

5.5.2 Further Reduction of S&P 500 Asset Numbers

For comparison, we first analysed the behaviour of the filters with the S&P 500 portfolio discussed in Chapter 4, when it was restricted to having the same number of assets as the Fx portfolio. Figure 5.3 shows the effect of RMT filtering on in-sample mean realised risk, for an S&P 500 portfolio consisting of 39 stocks, in the equally and exponentially weighted cases. These graphs are directly comparable with those for 100, 250 and 432 assets given in Section 4.4.4. The trend of the graphs from Section 4.4.4 was continued here, in the sense that the unfiltered forecasts became more competitive as the number of assets was reduced. We also noted, in the equally weighted case, that low Q -values were no longer optimal.

We found that filtering was effective for low numbers of stocks for certain values of the parameters. This was observed in the equally weighted case when the amount of historical data used was low, (equivalent to a low value of Q), and in the exponentially weighted case across a wide range of decay factors. However, the best unfiltered in-sample risk approached the best filtered result as the number of assets was reduced to 39.

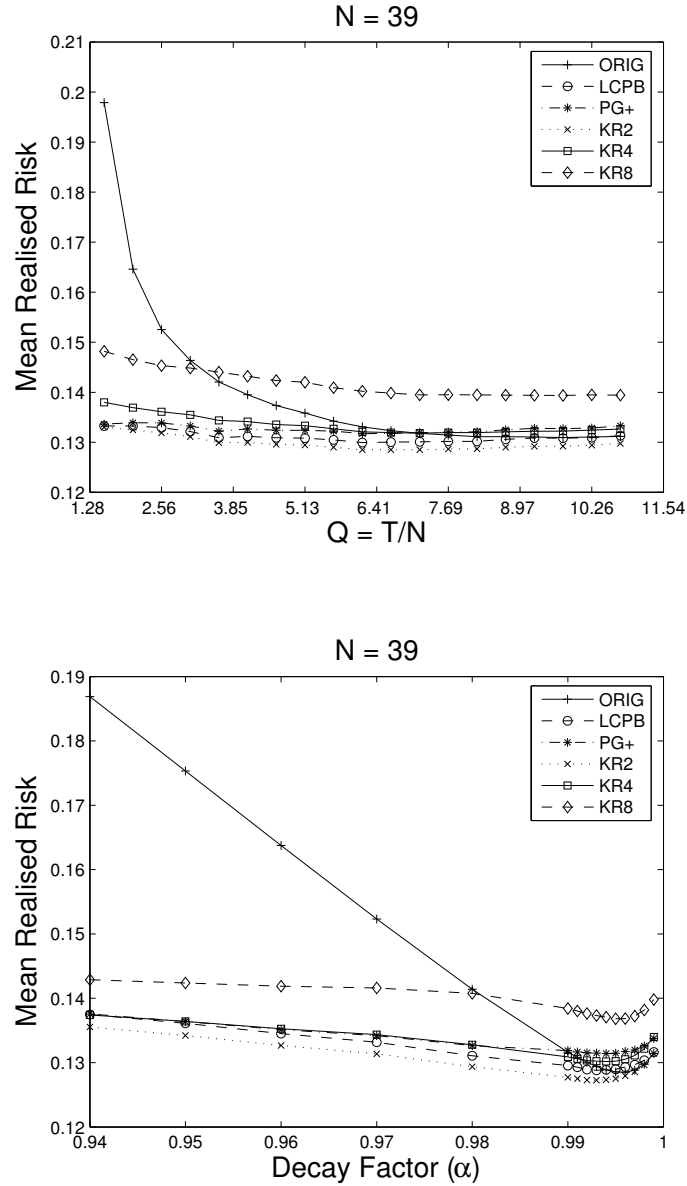


Figure 5.3: Mean bootstrapped (in-sample) realised risk, for selected filters, applied to equally (above) and exponentially (below) weighted volatility forecasts of S&P 500 stocks, and for unfiltered volatility (“ORIG”), for 39 assets. For equal weights, the x-axis shows the Q -value, which is the number of past moves used, divided by the number of assets in the portfolio. Note that the trend of the 432, 250 and 100 asset cases (Figures 4.5 to 4.9) was continued here, namely that the unfiltered forecasts became more competitive as the number of assets was reduced.

5.5.3 In-sample Analysis of the Fx Portfolio with 39 Assets

We now consider the in-sample results for the full Fx portfolio, with 39 assets. As in the S&P 500 case, the in-sample results showed the potential of RMT filters to reduce realised risk. Figure 5.4 shows the effect of RMT filtering on the in-sample mean realised risk of the Fx portfolio, with 39 assets, for equal and exponential weights.

Comparing these to the in-sample results for our S&P 500 portfolios we noticed a number of similarities. Taking the equally weighted case, we noted that, in general, the in-sample results for the Fx portfolio were very similar to those of the S&P 500 case. However, the type of improvements seen after filtering in the Fx case with 39 assets had more in common with the S&P 500 case with 432 assets, rather than the 39 stock case. Such similarities included the scale of risk reductions and the tendency for the lowest available number of past moves to be used (equivalent to the lowest Q value), in conjunction with RMT filtering.

In the exponentially weighted case, RMT filtering was again preferred overall. Here we noted a number of key differences to the stock portfolio. The unfiltered forecasts were more competitive in the Fx case, over a wide range of decay factors. Moreover, the optimal decay factor for these unfiltered forecasts was very close to the Riskmetrics (1996) value. For the filtered matrices, this optimal decay factor (in-sample) coincided exactly with the Riskmetrics (1996) recommendation of 0.97, and there was one clear global choice, instead of the local minima seen in the S&P 500 case.

In both the equally and exponentially weighted case, the KR2, KR4 and KR8 methods were again found to be amongst the best performing of all filters, and reasonably consistent with each other, while optimisation performance of the KR filters disimproved, in general, as the minimum replacement eigenvalue approached zero, meaning that, while some increase in stability was seen to improve risk performance, maximising stability was not effective.

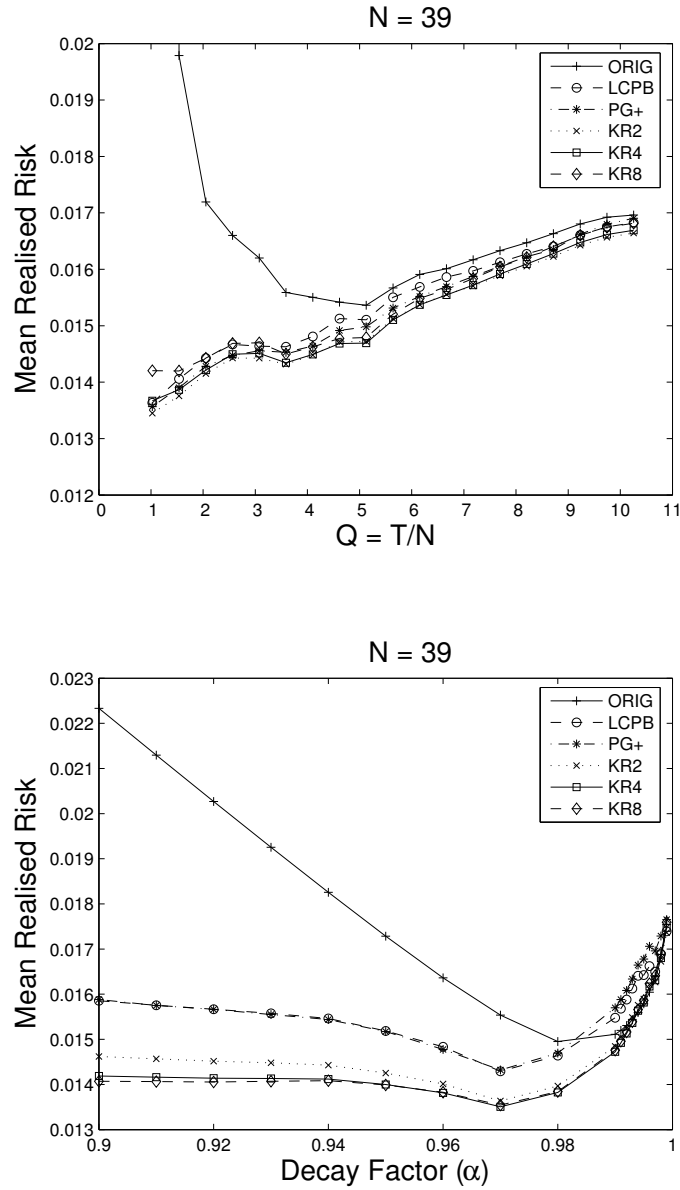


Figure 5.4: Mean bootstrapped (in-sample) realised risk, for selected filters, applied to equally (above) and exponentially (below) weighted volatility forecasts, and for unfiltered volatility (“ORIG”), for the Fx portfolio, with 39 assets. Filtering is seen to have been effective in both cases. For equal weights, a preference for low Q -values, with filtering, can be seen, and the closest agreement is with the S&P 500 case with $N = 432$, (Figure 4.6). In the exponential case, with filtering, the optimal decay factor (0.97) coincided with Riskmetrics (1996).

5.5.4 Further Reduction of Fx Asset Numbers

In this section we examine the effect that filtering had on an Fx portfolio with a reduced number of assets. Figures 5.5 to 5.8 show the effect of RMT filtering on in-sample mean realised risk, for portfolios with $N = 10, 15, 20$ and 30 assets, in the equally and exponentially weighted cases. These results shared many similarities with the results, for reducing the number of stocks in the S&P 500 portfolio, of Section 4.4.4.

In the case of equal weighting, it was observed that filtering was effective for a small Q value, at all four values of N , and so, in an environment where the number of past data points, T , was limited, either by choice or necessity, RMT filtering may be of some benefit. However, without restrictions on T , the effectiveness of RMT filtering was seen to be reduced as N was reduced.

For exponential weights, it was found that RMT filtering provided most benefit for lower decay factors, as in the S&P 500 case. However, in the presence of a free choice of decay factor, filtering was seen to provide little or no overall risk reduction for $N \leq 20$, as unfiltered estimates became more competitive with the reduction of assets, as was the case for the S&P 500 portfolio.

While the S&P 500 and Fx portfolios shared this tendency for filtering to be less effective as the number of assets was reduced, it was notable that this happened at different numbers of assets. Indeed, the value of filtering was already called into question for S&P 500 portfolios with 39 assets, while an Fx portfolio with the same number of assets clearly benefited from filtering. Meanwhile much lower asset numbers were involved before Fx portfolio filtering could be seen to be ineffective.

It is apparent from these results, and comparing the Fx and S&P 500 portfolios, that the relationship between the effectiveness of filtering, the number of assets, and the amount of information, is not as clear cut as previously concluded in the RMT literature.

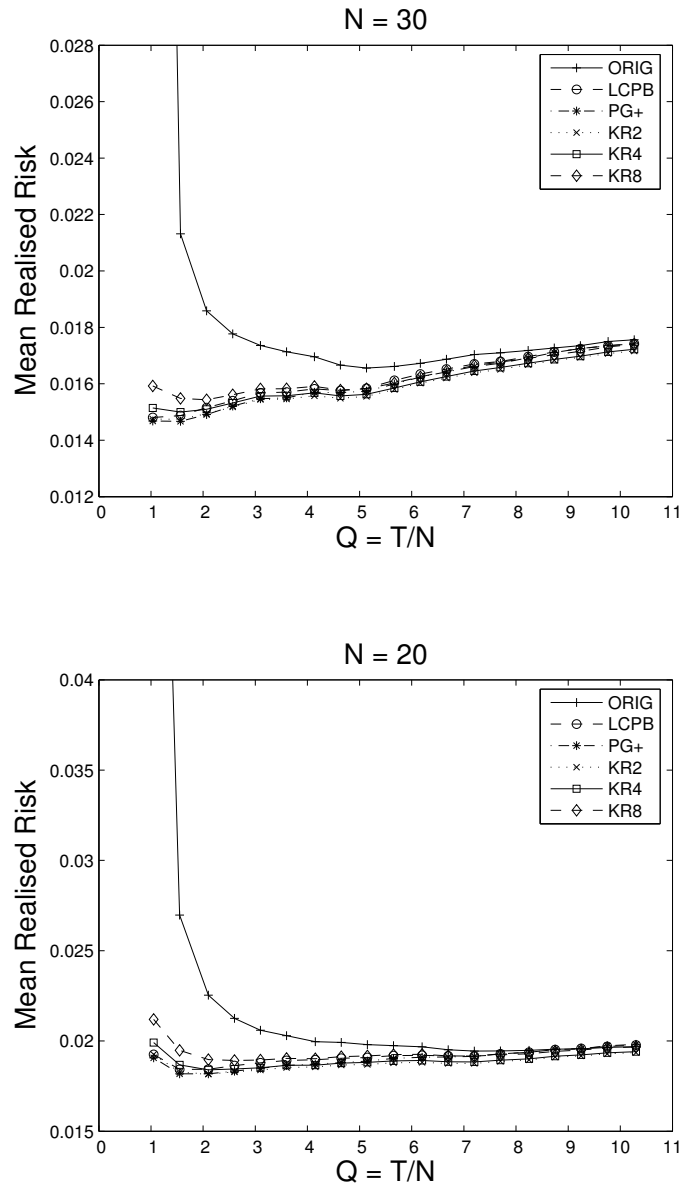


Figure 5.5: Mean bootstrapped (in-sample) realised risk, for selected filters, applied to equally weighted volatility forecasts of the Fx portfolio, and for unfiltered volatility (“ORIG”), for 30 (above) and 20 (below) assets. Filtering is still effective, reducing overall risk in both cases.

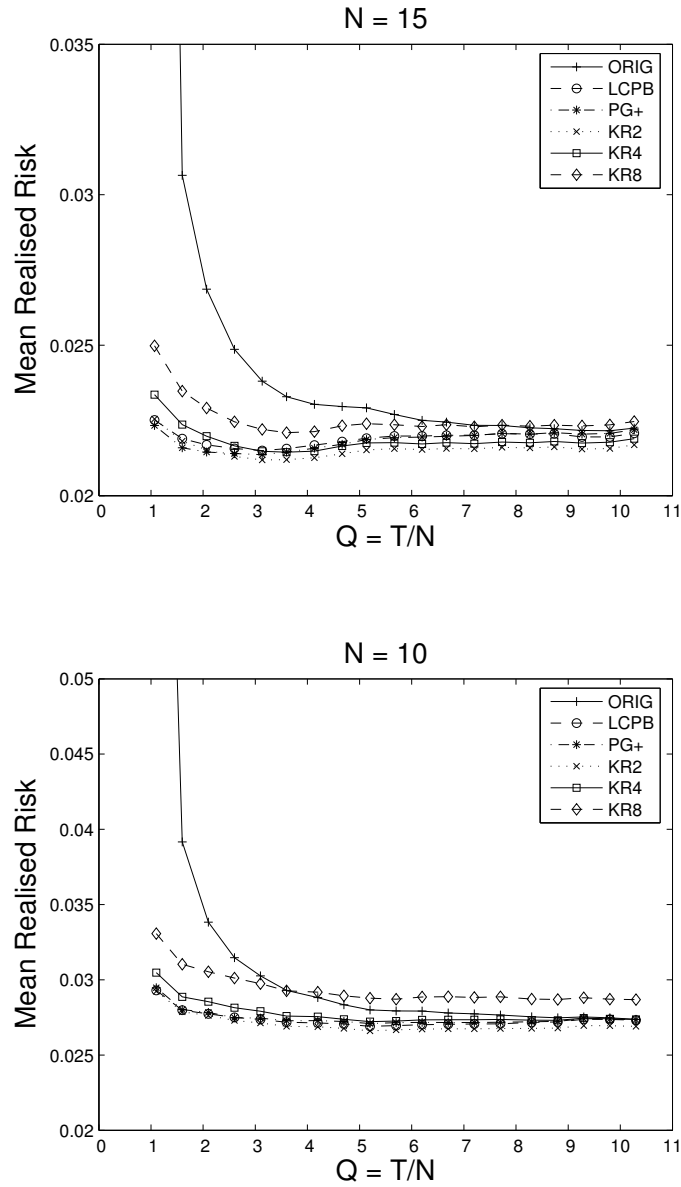


Figure 5.6: Mean bootstrapped (in-sample) realised risk, for selected filters, applied to equally weighted volatility forecasts of the Fx portfolio, and for unfiltered volatility (“ORIG”), for 15 (above) and 10 (below) assets. Note that as N was reduced, the unfiltered forecasts became more competitive.

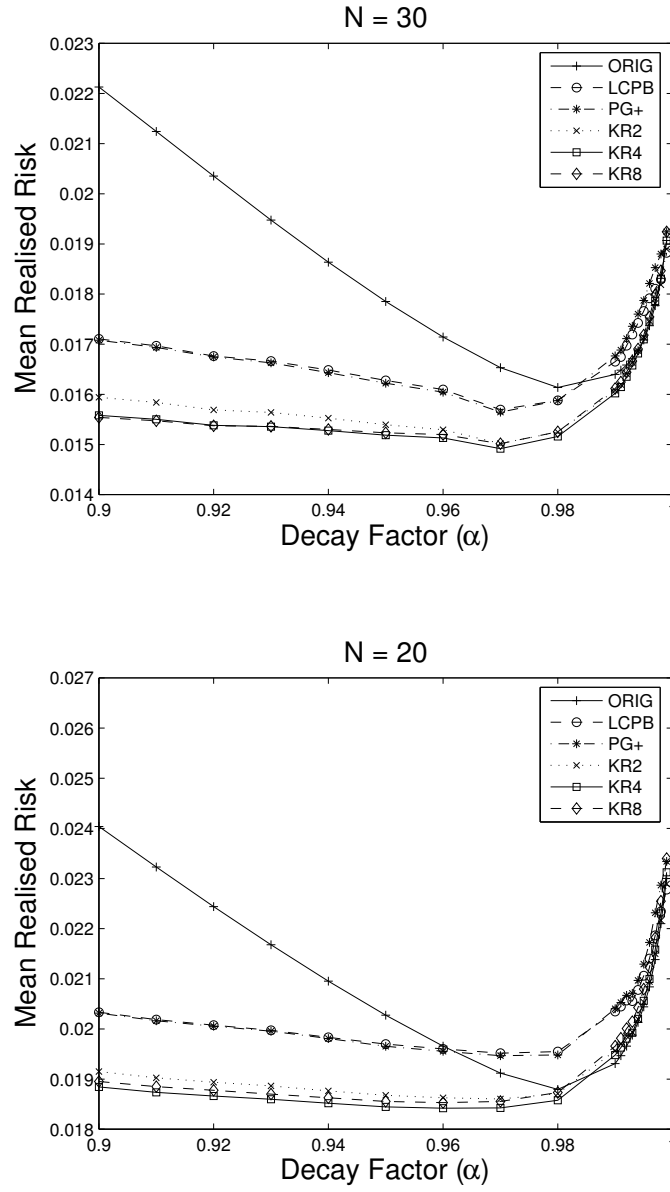


Figure 5.7: Mean bootstrapped (in-sample) realised risk, for selected filters, applied to exponentially weighted volatility forecasts of the Fx portfolio, and for unfiltered volatility (“ORIG”), for 30 (above) and 20 (below) assets. Stability filtering is still effective in both cases.

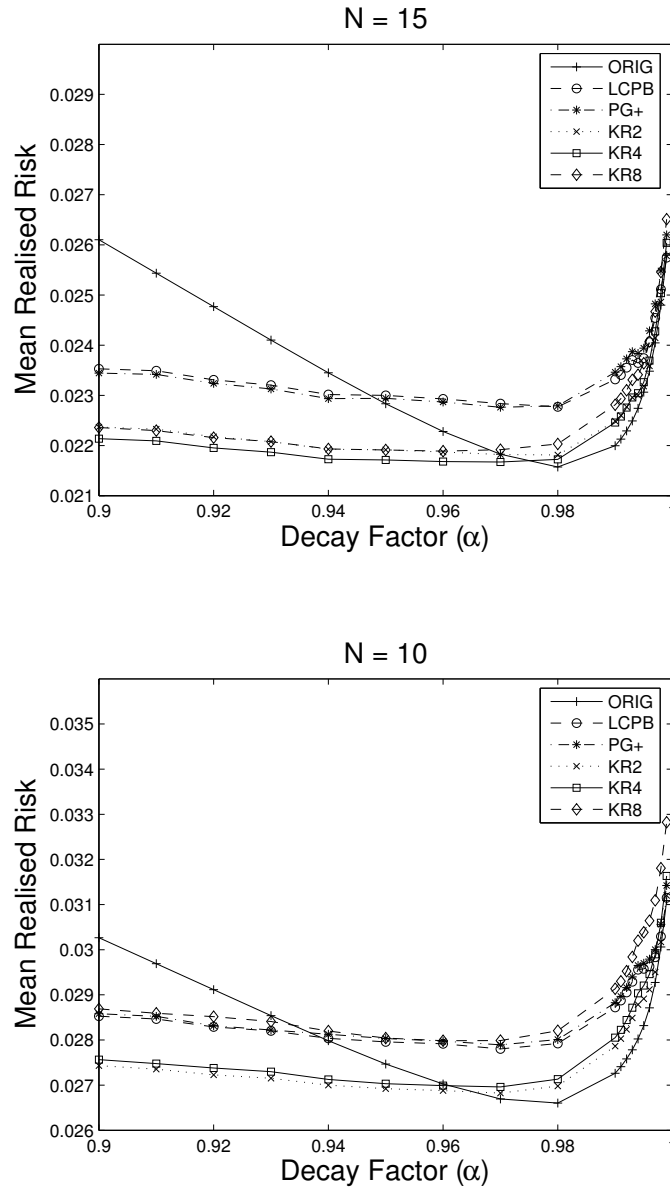


Figure 5.8: Mean bootstrapped (in-sample) realised risk, for selected filters, applied to exponentially weighted volatility forecasts of the Fx portfolio, and for unfiltered volatility (“ORIG”), for 15 (above) and 10 (below) assets. As for equal weights, the unfiltered forecast was shown to be more competitive as N was reduced.

5.6 Out-of-sample Testing

We now review the out-of-sample test results, first for the full Fx portfolio with 39 assets, and then for a sub-portfolio with 15 major assets.

5.6.1 Methodology

For the out-of-sample testing, we again compared the models using forward validation, as outlined in Section 3.3.2. To summarise, we considered every available test date and for each one used data prior to the test date to optimise any model parameters. The value of the weighting parameter (α or T) and the choice of KR model were determined out-of-sample.

In the Fx case, the forward validation was performed over a period of 1837 days, 129 of which were used as the initial training period, consistent with the S&P 500 analysis. Subsequent retraining was done daily. All 39 assets were used to eliminate the need to arbitrarily choose assets each day, unless otherwise stated.

5.6.2 Out-of-sample Analysis

5.6.2.1 Overall Risk

Table 5.2 shows a summary of the performance of the forecasting and filtering combinations. The figures shown are mean realised risk as a percentage of the result for unfiltered equally weighted covariance. RMT filtering was seen *on average* to reduce realised risk in all cases, compared to the unfiltered portfolio.

We found evidence that behaviour was consistent with the S&P 500 analysis in Chapter 4. Namely, risk was reduced in all cases and the reductions were of the same magnitude as the S&P 500 case. For the Fx portfolio the range of reduction was seen to be 13.0% to 14.2% for equal weights, and 5.5% to 10.9% for exponential (exponential figures are expressed as a percentage of the corresponding unfiltered exponential result). (These figures in the S&P 500 case were 12.4% to 14.9% for equal weights, and 5.7% to 10.1% for exponential.) Here, the best overall performance was seen when applying the PG+ filter to the equally

Table 5.2: Mean out-of-sample realised risk as a percentage of that for unfiltered equally weighted covariance, for the full Fx portfolio, with 39 assets. Filtering was seen to reduce mean realised risk in all cases.

Model	Unfiltered	LCPB	PG+	KR
Equal Weights	100	87	85.8	86.6
Exponential Weights	98.1	91.8	92.7	87.4

weighted forecast, while the KR method was also competitive when using equal weights. In the exponentially weighted case, KR filtering was found to be best, offering a 69.8% improvement in reduction compared to the best of the other models. We noted that in both the S&P 500 and Fx cases, the KR filter offered substantial improvements in reduction, for exponentially weighted forecasts, compared to the other filters.

5.6.2.2 Annual Risk

The annual mean realised risk of the different weightings and filters is shown, over 7 years, in Table 5.3. As in the S&P 500 case, we saw here a few instances where filtering increased the mean risk in a year. In this case, all filters were found to be capable of this behaviour. However, the majority of the time filtering did reduce risk. No very large risk increases were found in these annual figures. The largest percentage increase was 10.9%, when using exponential weights and the PG filter. When using equal weights this was reduced further, to 2.9%, in conjunction with the LCPB filter. In contrast, the best percentage risk reduction for equal weights was 31.7% with the LCPB filter, and for exponential weights was 23.5% with the KR filter. Therefore, the range of percentage changes after filtering were [-31.7%,2.9%] for equal weights, and [-23.5%,10.9%] for exponential (exponential figures are expressed as a percentage of the corresponding unfiltered exponential result). For the KR filter these ranges became [-31.2%,0.2%] for equal, and [-23.5%,4.9%] for exponential (exponential figures are expressed as a percentage of the corresponding unfiltered exponential result).

Table 5.3: Mean out-of-sample realised risk per year, for 7 years, as a percentage of the equally weighted unfiltered result each year, for the full Fx portfolio, with 39 assets. RMT filtering was seen to **increase** mean realised risk over the course of a year in a few cases, and for all filters, and to reduce it in the majority of cases.

Weights	Filter	1	2	3	4	5	6	7
Equal	None	100	100	100	100	100	100	100
Equal	LCPB	93.1	102.9	85.5	102.6	68.3	79.2	87
Equal	PG+	93.7	87.5	87.9	102.8	69.7	76.3	86.6
Equal	KR	95.2	100.2	88.8	99.3	68.8	76.2	87.4
Exp	None	95.8	115.5	104.3	102.7	90.8	88.4	103.2
Exp	LCPB	92.7	104.2	99.6	108.5	70.2	81.7	99.5
Exp	PG+	93	104.2	99.2	113.9	72.8	78.5	101.6
Exp	KR	88.5	91	90.5	107.7	69.5	76.9	96.1

5.6.2.3 Daily Risk

Table 5.4 shows the frequency of daily filtering effects. “Low” and “High” specify a range for the realised risk, expressed as a percentage of the relevant unfiltered realised risk. The effects on equally weighted and exponentially weighted matrices were combined here. In this case we found that, taking the mean across all filters, RMT filtering reduced realised risk on 62.6% of the days. The KR method was the most consistent in terms of reducing realised risk, doing so on 66.7% of the days, compared to 59.9% and 61.1% for the LCPB and PG+ methods respectively. We noted that these figures were reduced, compared to the corresponding ones for the larger S&P 500 portfolio, (which were reduction frequencies of 80.7% (KR), 71.7% (LCPB), and 70.6% (PG+) respectively). Combining all methods, RMT filtering of the Fx portfolio caused an increase in realised risk of 20% or more on 15.4% of the days. The stability-based KR filter had the lowest frequency of such increases (10.6%). This compared favourably with 16.9% for the LCPB method and 18.6% for the PG+ method. On 2.8% of the days these filters increased risk by 60% or more. This was somewhat higher than in the S&P 500 case (0.3%), while the size of the largest increase (and decrease) was also found to be greater than that of the S&P 500 case.

Table 5.4: Daily frequency, of percentage effect on realised risk, of applying RMT filters to the Fx portfolio. “Low” and “High” specify a range for the realised risk, expressed as a percentage of the relevant unfiltered realised risk. Similar to the S&P 500 case, RMT filtering was seen to reduce realised risk in the majority of cases, while also having the potential to increase realised risk, more than doubling it on some individual days.

Low	High	LCPB	PG+	KR
20	40	66	68	59
40	60	342	320	331
60	80	712	650	724
80	100	927	1050	1163
100	120	790	692	776
120	140	333	370	205
140	160	134	155	93
160	180	67	66	44
180	200	25	23	5
200	220	5	6	3
220	240	2	5	7
240	260	6	7	5
260	280	6	3	1
280	300	1	1	0

5.6.3 Parameter Values

Figure 5.9 shows the optimal number of past moves (in the equally weighted case), and decay factor values (in the exponential case), as selected through time by the forward validation. Here, unlike for the S&P 500, the decay factors chosen were quite consistent with the value of 0.97 suggested by Riskmetrics (1996), although we noted that the optimal value changed over the course of the test. For the Fx portfolio, the unfiltered decay factors were higher than the filtered ones, *but not, as was the case for the S&P 500, the maximum tested*. These results were consistent with the optimal in-sample Fx parameters.

In the equally weighted case, low numbers of past moves were preferred by the end, and so a tendency toward models which were reactive to recent market changes was seen. In fact, the best number of past moves in the equally weighted case tended toward the *lowest* possible value (40 days), indicating that using the previous two months data was optimal for forecasting the next month in this case. The best values for unfiltered forecasts were

higher, at 200 to 240 days. In both cases, this was consistent with the in-sample parameters.

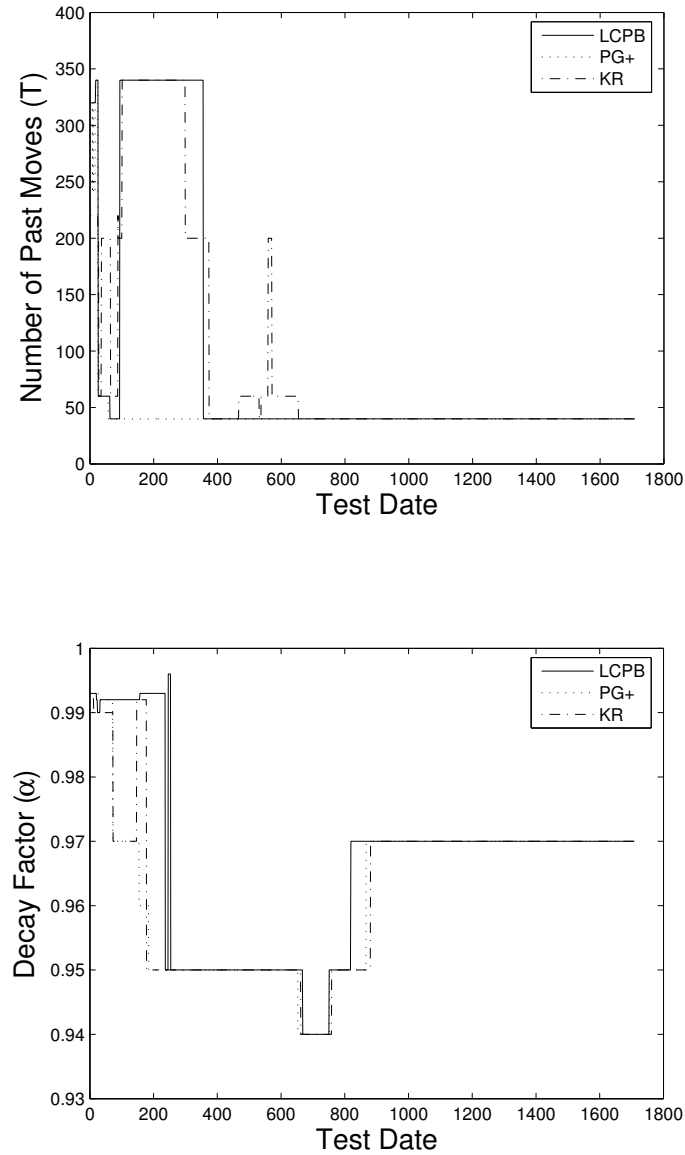


Figure 5.9: Optimal parameters for equally (above) and exponentially (below) weighted forecasts, as selected by forward validation, for the Fx portfolio. In the unfiltered cases (not shown), the number of past moves ranged between 200 and 240, and the decay factors were usually equal to 0.98, and sometimes in the range 0.99 - 0.995. A tendency toward reactive models can be seen in these graphs.

5.6.4 Range of Realised Risks

In Section 5.6.2 we saw that filtering could increase daily realised risk. Here, we study the effect of these increases on the *range* of the filtered risks, again finding that filter performance was not consistent across different risk levels. The filters performed better when unfiltered risk was high and so, as in the S&P 500 case, the range of the filtered risks was equivalent to that of the unfiltered risk. In the Fx case we found that risk was decreased on 62.6% of the days, while on 15.4% of the days risk was increased by 20% or more. *Here, both the frequency, and the size, of the large risk increases were bigger than for the S&P 500.* The largest percentage risk increases observed for Fx were in the range of 180% to 200%.

Figure 5.10 shows the effect of filtering, on each individual day, for equally and exponentially weighted covariance respectively. The unfiltered risk is included for comparison, as a straight line. As in the S&P 500 case, we found that the bulk of observations were at low unfiltered risk levels. We also observed good risk reduction when the unfiltered series was at the upper end of its range, while the largest percentage increases occurred in the middle and lower ends of the range. We again found that the *range* of the filtered risks was equivalent to that of the unfiltered risk. In fact we found that, for both weightings, the highest filtered risk was less than the highest corresponding unfiltered risk. Moreover, the overall frequency of high risks remained consistent, or reduced, after filtering.

These results are summarised in Tables 5.5 and 5.6, for equally and exponentially weighted covariance forecasts respectively. Here, as before, we broke the full range of unfiltered risks into quintiles. “All data” again refers to the full range, while “Top Quintile” refers to the top 20% of unfiltered risks, and so on. The results shown are the average risk, after filtering, in each quintile, as a percentage of the corresponding unfiltered result. Thus it can again be seen that filtering led to the biggest percentage risk reductions during periods where that risk was high. For equal weights we saw reductions of up to 28.4%, and, for exponential, 16.9%, in the top quintile. For all three filters, performance deteriorated with each subsequent lower quintile. For the lowest two quintiles we see risk was increased after

filtering in all cases for equal weights, and half the cases for exponential weights.

We also noted that, despite reducing the higher unfiltered risks, *none of the filters resolved the issue of poor forecasting performance in the upper end of the range*. An example of this is seen in Figure 5.11, which shows forecast and realised risks, for equally weighted covariance, comparing unfiltered and KR filtered results. The increased inaccuracy for forecasts in the upper part of the range can be clearly seen in both cases.

Table 5.5: Summary of filter performance in different unfiltered risk environments, for equally weighted covariance forecasts. “All data” refers to the full range of unfiltered realised risk. The remaining columns break this range into quintiles. The results shown are the average realised risk in each quintile, as a percentage of the corresponding unfiltered result. Thus it can be seen that the filters were most effective at reducing risk during periods where that risk was high.

Filter	All data	Top Quintile	Quintile 2	Quintile 3	Quintile 4	Bottom Quintile
None	100	100	100	100	100	100
LCPB	87	73.4	77.8	97.4	105	110.3
PG+	85.8	71.6	76.6	95.9	106.1	108.8
KR	86.6	75.8	77.2	96.8	102.6	103.8

Table 5.6: Summary of filter performance in different unfiltered risk environments, for exponentially weighted covariance forecasts. This table corresponds to Table 5.5. It can again be seen that the filters were most effective at reducing risk during periods where that risk was high.

Filter	All data	Top Quintile	Quintile 2	Quintile 3	Quintile 4	Bottom Quintile
None	100	100	100	100	100	100
LCPB	93.6	85.9	91.1	96.2	99.4	109.3
PG+	94.5	83.1	92.9	99.5	104.4	110.1
KR	89.1	85.5	87.5	90.6	92.7	95.5

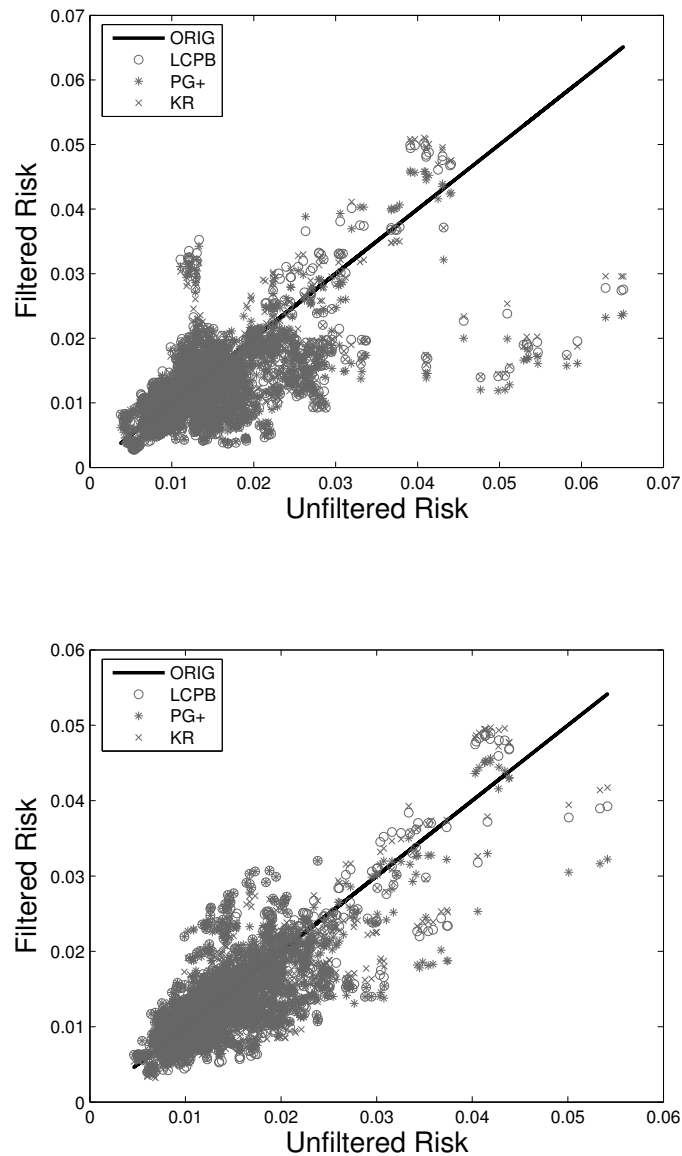


Figure 5.10: Filtered vs. unfiltered realised risks, for equally (above) and exponentially (below) weighted forecasts. The unfiltered risk is also included, as a 45 degree line, for comparison. It can be seen that filtering was most effective when the unfiltered series was in the upper end of its range, while the largest percentage risk increases occurred in the middle and lower ends of this range. Overall, the upper limit of the range of the filtered risks was less than that for the unfiltered risk in both cases.




Figure 5.11: Forecast and realised risks, for unfiltered (above) and KR filtered (below) equally weighted covariance forecasts, over the life of the forward validation. We see a shortfall in forecasting in the middle and, particularly, upper sections of the range. This “flat” forecast profile, and resulting shortfall, was common to all three filters, (and the unfiltered), and to both types of weighting scheme. This highlights the difficulty in forecasting financial markets.

5.6.5 Reduction of Assets

Following on from the in-sample analysis of Section 5.5.4, where the effect of filtering was seen to be reduced as the number of assets was decreased, we now review the out-of-sample behaviour of a portfolio consisting of 15 major currencies and commodities, selected from the main portfolio. Consistent with the in-sample results, RMT filtering was found to provide no benefit in the long run, despite some variation in the daily realised risk, in the range [-50%, +60%]. This is shown in Tables 5.7 to 5.9. Overall, filtering was seen to increase risk slightly in most cases, while risk was also increased for the majority of individual years and days.

Table 5.7: Mean out-of-sample realised risk as a percentage of that for unfiltered equally weighted covariance, for the Fx portfolio with 15 major assets. Risk was increased after filtering in all but one case.

Model	Unfiltered	LCPB	PG+	KR
Equal Weights	100	103.8	103.1	99.1
Exponential Weights	97.1	103.1	102.6	98.2

Table 5.8: Mean out-of-sample realised risk per year, for 7 years, as a percentage of the equally weighted unfiltered result each year, for the Fx portfolio with 15 major assets. Compared to the 39 assets case, we saw many more years with **increases** after filtering, and larger increases were observed.

Weights	Filter	1	2	3	4	5	6	7
Equal	None	100	100	100	100	100	100	100
Equal	LCPB	98.4	106.6	102	112.4	109.2	99.5	94
Equal	PG+	99	106.5	102.4	111.7	105.5	99.2	93.1
Equal	KR	96.2	101.8	99.4	103.2	99.3	96.2	96.5
Exp	None	93.5	98.7	99.5	98	93.3	95.9	103.6
Exp	LCPB	97.6	105.2	100.8	111	101.8	101	105.7
Exp	PG+	99	106.8	100.4	112.6	99.4	97.7	103
Exp	KR	95.9	101.3	99.2	102.1	94	94.6	102.3

Table 5.9: Daily frequency, of percentage effect on realised risk, of applying RMT filters to the Fx portfolio with 15 major assets. “Low” and “High” specify a range for the realised risk, expressed as a percentage of the relevant unfiltered realised risk. The range of frequencies was seen to be narrower than that of the full Fx Portfolio, while a higher frequency of risk increases was observed.

Low	High	LCPB	PG+	KR
50	60	2	3	0
60	70	18	7	13
70	80	72	45	25
80	90	328	399	182
90	100	856	882	1455
100	110	889	990	1441
110	120	704	635	267
120	130	358	308	32
130	140	161	107	1
140	150	21	34	0
150	160	7	6	0

5.7 Summary

In this chapter, we have studied the application of three RMT filters to a currency and commodity portfolio consisting of just 39 assets. We found that our results were in agreement with those for the S&P 500 portfolio, namely that RMT-based filtering can improve the realised risks of minimum risk portfolios, despite the *low number of assets* considered here.

Using forward validation, RMT filters were found, overall, to reduce mean realised risk in all cases tested, and in the majority of individual years. However, they were also found capable of increasing realised risk substantially on some individual days. While this latter observation is valuable, it would be unrealistic to expect any forecasting system to succeed every day, as markets adapt to the arrival of new information. We see that good annual consistency was observed, in the presence of these daily fluctuations. These out-of-sample results were consistent with the S&P 500 case.

We found that the days when risk was increased were more common when unfiltered risk was in the lower and middle parts of its range, and as a result, average risk was fre-

quently increased in the lower quintiles. We found good risk reduction when the unfiltered series was in the upper end of its range. We again found the biggest reductions in the top quintile of unfiltered risks, for both weightings, while performance was reduced with each lower quintile. Overall, this resulted in the *range* of the filtered risk being improved, compared to the unfiltered case, an important practical consideration.

When RMT filtering was applied to Fx portfolios with fewer asset numbers it was observed, in general, that the benefit of filtering was reduced as asset numbers decreased. In some cases filtering provided *no overall risk reduction*. This was also reflected in the out-of-sample filter performance, for a portfolio consisting of 15 major currencies and commodities, which revealed that, in this case, RMT filtering provided no long term risk reduction, and was more likely to *increase* realised risk, both overall and on any individual day.

When comparing the Fx results to those for the S&P 500, we noted that this loss of filter effectiveness occurred at different numbers of assets in the two markets, suggesting that it was reducing the number of assets from that of the “full” portfolio that caused the fall off in filter performance. The type of improvements seen after filtering in the Fx case with 39 assets had more in common with the full S&P 500 case with 432 assets, rather than the 39 stock case. Such similarities included the scale of risk reductions and the tendency for the lowest available number of past moves to be used with RMT filtering. Therefore, the relationship between the effectiveness of filtering, the number of assets, and the amount of information, is less straightforward than previously discussed in the literature (Pafka and Kondor, 2002a,b). This introduces the possibility of applying RMT filters to a wider range of markets, namely those with a high ratio of data to assets. We consider the power of past data for forecasting, and the complex interactions of the system, rather than simply the ratio of available data to assets, to be important in this context, based on these results.

RMT filters also uncovered different styles of models than were possible with unfiltered analysis, namely ones that reacted quickly to market events. Without filtering these models utilising very recent data were found to be hidden by noise. Meanwhile, *in contrast to the* S&P 500 case, the decay factors chosen here showed good consistency with the value of

0.97 suggested by Riskmetrics (1996).

The observed behaviour of the stability-based filter was generally in agreement with that of the S&P 500 case, namely that the KR2, KR4 and KR8 methods delivered better performance than most filters, and were reasonably consistent with each other. When considering exponential weights, the KR filter performed best over the period, and was the best filter in each individual year. In total, it improved risk reduction by 69.8% for exponential weights. It also remained competitive when using equal weights, although in that case the PG+ filter was marginally preferred. The KR filter was also found to reduce the number of days that filtering led to a large increase in risk, as well being the filter that reduced risk most often, as in the S&P 500 case. Here, the KR filter reduced risk on 66.7% of the days, compared to 59.9% and 61.1% for LCPB and PG+. Finally, the KR filter again performed consistently well across the full range of unfiltered risks during this test, and was again the only filter to reduce in each quintile for exponential weights.

Taken as a whole, these results suggested that RMT filtering can provide strong risk reduction, in this case of a foreign exchange and commodity portfolio. Filtering was found to reduce average portfolio risk by between 5.5% and 14.2%, and to do so in a way that was consistent with a larger portfolio.

Chapter 6

Conclusions and Future Work

6.1 Goals of the Thesis

Here, we summarise the goals of the thesis, which were

- an assessment of the stability-based KR filter. We wished to study this, and the other two filters mainly from the point of view of risk reduction, but also taking stability and consistency into account.
- improving the test methodology, (c.f. that reported in previous literature), by including in particular an out-of-sample test, which simulated a live implementation where parameters were updated, and trades were executed, on a daily basis.
- the application of the filters to a medium term forecasting scenario, while making minimal assumptions about the trading strategy.
- a review of the potential use of RMT filters in the foreign exchange market.
- an examination of filter behaviour on an annual, monthly and daily basis, to complement the summary results found in the literature, and here.
- studying the choice of forecast weighting, and the best parametrisation of these methods, with and without filtering.

- measuring the impact of filtering the covariance matrix directly, as opposed to filtering correlation.

6.2 Conclusions

In this section we review and discuss the conclusions of the thesis, from the point of view of a practitioner considering implementing one of these forecasting systems.

6.2.1 Overview

In this thesis we have examined the application of three RMT filters, including one novel filter, to the optimisation of (1) an S&P 500 and (2) a foreign exchange portfolio. We have studied the effect on realised portfolio risk, both in and out-of-sample. Our results are in good agreement with previous results (Pafka et al., 2004), that RMT-based filtering can improve the realised risk of minimum risk portfolios. Similar in-sample reductions were seen for both portfolios, and corresponded well with Pafka et al. (2004).

We observed that, for the full portfolios, filtering correlation, these three filters reduced risk out-of-sample in all cases, with overall improvements ranging from 5.5% to 14.9%. We primarily used a forecasting period of 20 days for our analysis, while results for a 50 day forecast were found to agree well with the shorter time frame.

When considering the effect of filtering on annual risk, good consistency was observed. Overall we tested correlation filtering over 13 years, spread between the two asset classes. This was done for both equal and exponential weighting and for all three filters. This resulted in 78 annual tests of RMT filtering. Of these, risk was reduced 88.5% of the years. The best annual reduction measured was 31.7%.

When the daily figures were examined, we found risk reduction in the majority of portfolio trades: 74.3% and 62.6% of the days for the stock and foreign exchange portfolios respectively. The fact that risk was capable of being increased, and moreover in so many cases, was to our knowledge an effect that had not been previously discussed in the litera-

ture, and one that would be vital knowledge to anyone planning to implement such filters. In addition, all of these filters were capable of substantial risk increases. These were measured at up to 200% for Fx, and 100% for the stocks.

While these limitations should be noted, it would be unrealistic to expect that any forecasting system be 100% accurate, due to the arrival of new market information. Despite these daily fluctuations, consistency *was* achieved annually. Moreover, the *range* of the filtered risks was found to be equivalent, or slightly improved on, that produced in the absence of filtering. This resulted from the large increases occurring when the unfiltered risk was already low. On the days when the unfiltered was high, filtering performed well.

6.2.2 Stability Filtering

In this thesis, we have proposed a novel method of filtering covariance matrix forecasts with RMT. This new filter is based on improving the stability of the filtered matrix, to prevent optimised solutions being too sensitive to changes in the underlying market. While maximising this stability preferentially could not be recommended, due to its adverse effect on realised risk, *balance between stability and risk performance* is vital to a successful outcome, and in this case, when filtering correlation, led to

- improved out-of-sample risk reduction, by 49.2% for the stock portfolio, and 69.8% for the foreign exchange portfolio¹, in the case of exponential weights.
- a moderate improvement (17.3%) to risk reduction in the case of equal weights, and the stock portfolio. The stability filter was also competitive in the equally weighted case for foreign exchange.
- increased Krzanowski stability (c.f. Section 3.2.6.1) compared to established methods. (This is an expected result due to the definition of the filter, which is designed to increase Krzanowski stability)
- an increase in the percentage of days that risk was reduced.

¹These improvements were relative to the best performing alternative

- a decrease in the number of days where filtering resulted in large risk increases.
- good consistency of risk reduction, both on an annual basis, and at different levels of the unfiltered risk.

The frequency of daily reduction was measured at 80.7% for the stocks, while that for foreign exchange was 66.7%. We found a good agreement between the most competitive KR filters, namely the KR2, KR4 and KR8 filters. These involved minimum replacement eigenvalues of one half, one quarter and one eighth of the average noisy eigenvalue. For the stock portfolio, the KR filter was the *only one to reduce each year*, and only once did another filter outperform it, while for foreign exchange, it was the *best filter in every year for exponential weights*. Meanwhile, for exponential weights, and in both markets, the stability filter was the only one to reduce in all five quintiles of the unfiltered realised risk.

6.2.3 Foreign Exchange Filtering

Despite the fact that RMT filtering was designed for situations where the ratio of data to assets was low, we found that these filters applied well to an Fx portfolio with just 39 assets, and a long time series of data. The effects of filtering, both in and out-of-sample were very similar to an S&P 500 portfolio with 432 assets. In the in-sample case, and with equal weights, the similarity was striking. In both cases, using too little price history without filtering resulted in large risks, while using too much was also not ideal. When filtering was added, the picture changed completely for both asset classes, and in the same way. Now, using the shortest possible price history was best. In the Fx case this was 40 data points, just two months of data. This is a very reactive model, and indicates that it is not the *availability* of past data that is important in this case, but the *relevance* of that data for forecasting.

In the exponential case, similarities were also seen between the two portfolios in-sample. In the unfiltered exponential case, a wide range of decay factors performed poorly, particularly at lower values. With filtering included, risk was reduced substantially across this range in both cases.

Table 6.1: Mean realised risk as a percentage of unfiltered, for all filters and both portfolios. The improvements show remarkable consistency between the two asset classes.

Weights	Filter	Fx	S&P 500
Equal	LCPB	87	87.3
Exp	LCPB	93.6	93.3
Equal	PG+	85.8	87.6
Exp	PG+	94.5	94.3
Equal	KR	86.6	85.1
Exp	KR	89.1	89.9

Out-of-sample, the results were also very similar. The improvements seen are shown in Table 6.1. Here, we see just how similar the filter behaviour is between these two portfolios, giving a clear impression that filtering can be applied successfully to this smaller portfolio.

The similarities also extended to the yearly and daily figures. Here, we saw risk reduction in the majority of years, with some years showing moderate increases. Daily ranges were wider for the Fx portfolio, and the frequency of reduction was lower. However, the majority of days involved risk improvement for both markets, and the histograms showed similar patterns. In both cases, the best reduction occurred when the unfiltered value was high and so the range of realised risk was effectively unchanged, or improved, by filtering.

All of these results point to RMT filtering being very valuable as a tool for improving the risk of foreign exchange investments. However, filter use should be taken into consideration in conjunction with choosing the composition of assets in the portfolio. Here, we have included all currencies that were tradeable. Any attempt to reduce the number of currencies resulted in a fall off in filter performance. A similar effect was seen for stocks, although at different asset numbers. These results suggest that filtering will be more effective for managers who trade a more diversified portfolio, and that those who restrict their asset choice should reconsider this strategy, in light of this new technology.

6.2.4 Choice of Weighting System and Parametrisation

We now consider the choice of weighting system in the forecasts, and the associated parameters. Without filtering, exponential weights outperformed equal weights in all out-of-sample tests. However, it has been noted that the decay factors that were associated with this improvement were far from those recommended by Riskmetrics. In many cases, the optimal unfiltered decay factors approached one, and therefore lost much of the intention behind using exponential weights, performing almost like an equally weighted forecast instead.

When filtering was included, we saw mixed results. As is clear from Table 6.1, the effect of filtering was always greater for equal weights. As a result, in the S&P 500 case, the best model was exponentially weighted, while in the Fx case equal weights performed better. When using the stability based filter, there was little to choose between the different weightings. Again, exponential weights were better for the S&P 500 portfolio and equal for the Fx one.

We recall that in Riskmetrics (1996), it was considered preferable to use exponential weights, due to conceptual benefits. In particular, the ability of events to suddenly drop off the end of the equal weighting moving window, and the fact that recent changes in volatility are given the same weighting as older information, were considered drawbacks. From our results, when noise filtering was included, the difference in performance of the two weightings was not great, and in some cases, equal weights were better. This suggests that those times when giving a current event an equal weighting is a disadvantage, are compensated by other times when a sudden shift in market behaviour turns out not to be sustained, and the longer term behaviour is soon restored.

In Pafka and Kondor (2002a,b), it was discussed that the effect of noise was reduced by increasing the amount of data being used. From our results it became clear that this was not the complete picture. While the unfiltered and filtered bootstrap results did indeed converge as the amount of data increased, we saw that this convergence was not to the best risk. Instead, by reducing the amount of data, in the equally weighted case, to the minimum,

and filtering, a much lower risk was found. Therefore, ideas based on simply increasing the amount of data, to remove the need for filtering, are overly simplistic.

This improvement, (after filtering), in the reaction times to market events, was seen in both markets, and also reflected in the optimal decay factors. We saw that these decay factors were reduced, in general, after filtering, particularly for the stock portfolio. The optimal out-of-sample decay factors for the S&P 500 were still higher, after filtering, than the Riskmetrics value, while in the Fx case the filtered decay factors showed good agreement with that work.

6.2.5 Covariance Filtering

We have also considered direct filtering of the covariance matrix, since it contains more information about the individual assets, and was also shown to contain more non-noisy eigenvalues. In this case the S&P 500 portfolio was used, and we found that, while direct filtering reduced risk in all out-of-sample tests, filtering the correlation matrix produced lower risk. In fact, filtering correlation resulted in lower risk in five of the six test cases. The frequency of reductions also disimproved when covariance was filtered, while larger increases were observed. Finally, the overall best forecast involved correlation filtering. Based on these tests, applying RMT filtering to the correlation matrix was preferred.

6.2.6 Summary

In these tests it was seen that the use of RMT filtering reduced the realised risk of investment portfolios. All filters were successful. In the case of exponential weights, our novel stability filter out-performed the other filters tested. The situation was more competitive for equal weights, where the stability filter was best for the stock portfolio. In general, this filter showed greater consistency throughout.

Despite the small number of assets involved, filtering was just as valid for foreign exchange as for stocks, implying that noise cannot be reduced simply by considering more data. In fact, for both markets, filtering led to models which considered less data, and were

thus more reactive to recent events in the market.

6.3 Future Work

Following on from this, we now consider future work. First, since assumptions have been made here, to preserve the generality of the test to many different applications, our next step would be to test one specific implementation.

In a hedge fund environment the main measure of success is the ratio of the trading return to trading risk, known as the Sharpe ratio (e.g. L'Habitant (2006)). Therefore, any system which uses filtering to reduce risk can gain a competitive advantage. Moreover, funds are allocated to managers not only on their own Sharpe ratio, but also on their correlation to other managers, and so there is a definite advantage, in a portfolio of funds context, associated with reducing correlation to other managers. Any new technology is beneficial in this regard. With this in mind, we consider the typical steps necessary for developing this technology into a working trading system. This would involve

- a forecasting system for expected returns. These forecasts are critical to the success of any systematic trading model. With forecasted returns available, we can consider the problem of minimising risk for a fixed target level of return. The literature (e.g. Laloux et al. (2000); Conlon et al. (2007)) records a steady level of risk reduction, after filtering, along the whole of the efficient frontier, at realistic target returns. It follows that comparable reduction to that seen in this thesis can be expected for any practical target return level.
- considered operating constraints. Many funds involve gearing² of their capital, which is possible in this case due to only a percentage of capital, known as margin, being needed at any one time to insure against trading losses. This, and other operating constraints, may be beneficial. However, recent work (e.g. Galluccio et al. (1998); Gábor and Kondor (1999); Kondor (2000); Bongini et al. (2002)) has pointed out that

²i.e. investment based on notional, or borrowed, funds. Gearing therefore increases potential profit and loss.

the addition of such non-linear constraints causes the appearance of many different unstable optimal solutions. Further research is clearly needed in this area.

- an accurate trading simulation system. Great care is needed in the building of such a system to avoid unrealistic predictions. A key element of this is a good understanding, and modelling, of spread costs and brokerage fees that are encountered in practice.

These are the key modelling questions. Following these, technological challenges such as database construction, trade execution and administration, reporting, and compliance, can be considered.

The second strand of future work relating to this thesis involves the improvement of the filters themselves. Second order improvements can potentially be made, by making adjustments to take account of the largest measured eigenvalues. Further details of this can be found in, e.g., Laloux et al (2000). Essentially, the total risk in the system can be iteratively adjusted, when calculating the maximum random eigenvalue, to exclude the effect of the eigenvalues known to be non-random. In this way further non-random sub-dominant eigenvalues may potentially be identified. Moreover, adjustments can also be made for finite N effects (Laloux et al., 2000). What the practical benefits of these refinements are, and whether they would be enough to compensate for the increased run time, is a matter for further research.

Finally, it would be interesting to study the correspondence between RMT and traditional models of mathematical finance, such as the implications of simulating data using RMT in combination with the Capital Asset Pricing Model or Arbitrage Pricing Theory (e.g. Campbell et al. (1996)).

References

- Bak, P., Paczuski, M., and Shubik, M. (1997). Price variations in a stock market with many agents. *Physica A*, 246:430–453.
- Bongini, L., Esposti, M. D., Giardina, C., and Schianchi, A. (2002). Portfolio optimization with short-selling and spin-glass. *Eur. Phys. J. B*, 27:263–272.
- Bouchaud, J. and Potters, M. (2000). *Theory of Financial Risks - From Statistical Physics to Risk Management*. Cambridge University Press, U.K.
- Bouchaud, J. and Potters, M. (2003). *Theory of Financial Risk and Derivative Pricing- From Statistical Physics to Risk Management*. Cambridge University Press, U.K.
- Brigo, D. and Hanson, B. (1998). On some filtering problems arising in mathematical finance. *Insur. Math. Econ.*, 22(1):53–64.
- Campbell, J. Y., Lo, A. W., and MacKinlay, A. C. (1996). *The Econometrics of Financial Markets*. Princeton University Press, U.S.
- Conlon, T., Ruskin, H. J., and Crane, M. (2007). Random matrix theory and fund of funds portfolio optimisation. *Physica A*, 382(2):565–576.
- Dacorogna, M., Gençay, R., Muller, U., Olsen, R., and Pictet, O. (2001). *An Introduction to High-Frequency Finance*. Academic Press, U.S.
- Daly, J., Crane, M., and Ruskin, H. J. (2008). Random matrix theory filters in portfolio optimisation: A stability and risk assessment. *Physica A*, 387(16):4248–4260.

- Daly, J., Crane, M., and Ruskin, H. J. (2009). Random matrix theory filters and currency portfolio optimisation. *Forthcoming in J. Phys. Conf. Ser. as part of the proceedings of APFA7: Applications of Physics in Financial Analysis 7, Tokyo, March 2009.*
- Dyson, F. (1962). Statistical theory of the energy levels of complex systems i, ii and iii. *J. Math. Phys.*, 3(1):140–175.
- Dyson, F. and Mehta, M. (1963). Statistical theory of the energy levels of complex systems iv. *J. Math. Phys.*, 4(5):701–712.
- Elton, E. J., Gruber, M. J., Brown, S. J., and Goetzmann, W. N. (2006). *Modern Portfolio Theory and Investment Analysis*. Wiley, U.S.
- Gábor, A. and Kondor, I. (1999). Portfolios with nonlinear constraints and spin glasses. *Physica A*, 274:222–228.
- Galluccio, S., Bouchaud, J., and Potters, M. (1998). Rational decisions, random matrices and spin glasses. *Physica A*, 259:449–456.
- Gençay, R., Selçuk, F., and Whitcher, B. (2002). *An Introduction to Wavelets and Other Filtering Methods in Finance and Economics*. Academic Press, U.S.
- Heath, A., Upper, C., Gallardo, P., Mesny, P., and Mallo, C. (2007). Bank for International Settlements, triennial central bank survey, December 2007, foreign exchange and derivatives market activity in 2007. Technical report, <http://www.bis.org/publ/rpfx07t.pdf>.
- Hull, J. (2006). *Options, Futures, and Other Derivatives*. Prentice Hall, U.S.
- Keogh, G., Sharifi, S., Ruskin, H. J., and Crane, M. (2003). Epochs in market sector index data - empirical or optimistic? In *Proceedings of the 2nd Nikkei Econophysics Symposium - Application of Econophysics, H. Takayasu (Ed.)*.
- Kondor, I. (2000). Spin glasses in the trading book. *Int. J. Theor. Appl. Finance*, 3(3):537–540.

- Krzanowski, W. (1984). Sensitivity of principal components. *J. Royal Stat. Soc. B*, 46(3):558–563.
- Kwapień, J., Drożdż, S., and Speth, J. (2004). Time scales involved in emergent market coherence. *Physica A*, 337:231–242.
- Laloux, L., Cizeau, P., Bouchaud, J., and Potters, M. (1999). Noise dressing of financial correlation matrices. *Phys. Rev. Lett.*, 83(7):1467–1470.
- Laloux, L., Cizeau, P., Potters, M., and Bouchaud, J. (2000). Random matrix theory and financial correlations. *Int. J. Theor. Appl. Finance*, 3(3):391–397.
- L’Habitant, F. S. (2006). *Handbook of Hedge Funds*. Wiley, U.K.
- Litterman, R. and Winkelmann, K. (1998). Estimating covariance matrices. Technical report, Goldman Sachs, Risk Management Series.
- Mantegna, R. and Stanley, H. (2000). *An Introduction to Econophysics*. Cambridge University Press, U.K.
- Markowitz, H. (1959). *Portfolio Selection: Efficient Diversification of Investments*. Wiley, New York.
- Mehta, M. (1991). *Random Matrices*. Academic Press, New York.
- Mehta, M. and Dyson, F. (1963). Statistical theory of the energy levels of complex systems v. *J. Math. Phys.*, 4(5):713–719.
- Morgan, J. P. and Reuters (1996). Riskmetrics technical document. Technical report, <http://www.riskmetrics.com/techdoc.html>.
- Nilantha, K., Ranasinghe, and Malmini, P. (2007). Eigenvalue density of cross-correlations in Sri Lankan financial market. *Physica A*, 378(2):345–356.
- Pafka, S. and Kondor, I. (2002a). Noisy covariance matrices and portfolio optimization. *Eur. Phys. J. B*, 27(2):277280.

- Pafka, S. and Kondor, I. (2002b). Noisy covariance matrices and portfolio optimization ii. *Physica A*, 319:487–494.
- Pafka, S., Potters, M., and Kondor, I. (2004). Exponential weighting and random-matrix-theory-based filtering of financial covariance matrices for portfolio optimization. *To appear in Quant. Fin.; e-print cond-mat/0402573*.
- Plerou, V., Gopikrishnan, P., Rosenow, B., Amaral, L., Guhr, T., and Stanley, H. (2002). Random matrix approach to cross correlations in financial data. *Phys. Rev. E*, 65(6):066126.
- Plerou, V., Gopikrishnan, P., Rosenow, B., Amaral, L., and Stanley, H. (1999). Universal and nonuniversal properties of cross correlations in financial time series. *Phys. Rev. Lett.*, 83(7):1471–1474.
- Plerou, V., Gopikrishnan, P., Rosenow, B., Amaral, L., and Stanley, H. (2000a). Econophysics: financial time series from a statistical physics point of view. *Physica A*, 279:443–456.
- Plerou, V., Gopikrishnan, P., Rosenow, B., Amaral, L., and Stanley, H. (2000b). A random matrix theory approach to financial cross-correlations. *Physica A*, 287:374–382.
- Plerou, V., Gopikrishnan, P., Rosenow, B., Amaral, L., and Stanley, H. (2001). Collective behaviour of stock price movements - a random matrix theory approach. *Physica A*, 299:175–180.
- Potters, M., Bouchaud, J., and Laloux, L. (2005). Financial applications of random matrix theory: Old laces and new pieces. In *Proceedings of the Kraków conference on Applications of Random Matrix Theory to Economy and Other Complex Systems*, *arXiv:physics/0507111*.
- Sengupta, A. and Mitra, P. P. (1999). Distributions of singular values for some random matrices. *Phys. Rev. E*, 60(3):3389–3392.

- Sharifi, S., Crane, M., Shamaie, A., and Ruskin, H. J. (2004). Random matrix theory for portfolio optimization: a stability approach. *Physica A*, 335:629–643.
- Sharkasi, A., Crane, M., and Ruskin, H. J. (2006a). Apples and oranges: the difference between the reaction of the emerging and mature markets to crashes. In *Proceedings of the Third Nikkei Econophysics Symposium - Practical Fruits of Econophysics*.
- Sharkasi, A., Crane, M., Ruskin, H. J., and Matos, J. A. (2006b). The reaction of stock markets to crashes and events: A comparison study between emerging and mature markets using wavelet transforms. *Physica A*, 368(2):511–521.
- Sornette, D. and Johansen, A. (1998). A hierarchical model of financial crashes. *Physica A*, 261:581–598.
- Strang, G. (1980). *Linear Algebra and its Applications*. Academic Press, New York.
- Tsay, R. S. (2005). *Analysis of Financial Time Series*. Wiley, U.S.
- Tumminello, M., Aste, T., DiMatteo, T., and Mantegna, R. N. (2005). A tool for filtering information in complex systems. *Proc. Natl. Acad. Sci. USA*, 102(30):10421–10426.
- Utsugi, A., Ino, K., and Oshikawa, M. (2004). Random matrix theory analysis of cross correlations in financial markets. *Phys. Rev. E*, 70(2):026110.
- Wilcox, D. and Gebbie, T. (2007). An analysis of cross-correlations in an emerging market. *Physica A*, 375(2):584–598.
- Zhu, M., Chiarella, C., He, X., and Wang, D. (2009). Does the market maker stabilize the market? *Physica A*, 388:3164–3180.

Appendix A

Maximum Eigenvalue of an Exponentially Weighted Random Matrix

A.1 Introduction

Many of the tests performed as part of this thesis involved significant runtime, and it was necessary to improve on this at various points. In this appendix we present an efficient method for determining the maximum eigenvalue of an exponentially weighted random matrix. This has been developed following from the work of Pafka et al (2004), who outlined the original formulation, as described in Section 3.2.4.

A.2 Maximum Eigenvalue of an Exponentially Weighted Random Matrix

As discussed in Section 3.2.4, Pafka et al. (2004) have shown that for a matrix $\mathbf{M} = \{m_{ij}\}_{i,j=1}^N$ of the form

$$m_{ij} = \sum_{k=0}^{\infty} (1 - \alpha) \alpha^k x_{ik} x_{jk} \quad (\text{A.1})$$

where $\{x_{ik}\}_{k=0,\dots,\infty}^{i=1,\dots,N}$ are $N.I.D.(0, \sigma^2)$, the special case, with $N \rightarrow \infty$, $\alpha \rightarrow 1$ and Q fixed, where

$$Q = \frac{1}{N(1 - \alpha)} \quad (\text{A.2})$$

results in the density, $\rho(\lambda)$, of the eigenvalues of \mathbf{M} being given by

$$\rho(\lambda) = \frac{Qv}{\pi} \quad (\text{A.3})$$

where v is the root of

$$F(v) = \frac{\lambda}{\sigma^2} - \frac{v\lambda}{\tan(v\lambda)} + \ln(v\sigma^2) - \ln(\sin(v\lambda)) - \frac{1}{Q} \quad (\text{A.4})$$

$F(v)$ is well defined on the open interval $(0, \pi/\lambda)$. If a root does not exist on this interval for a given value of λ we define $\rho(\lambda) = 0$ for that λ .

We now analyse the behaviour of this function, $F(v)$, on $(0, \pi/\lambda)$, and determine a more efficient method for finding the maximum eigenvalue. While this maximum eigenvalue can be found directly from Equation (A.4), that method involved a large computational expense, during our tests.

We first note that, on the interval $v \in (0, \pi/\lambda)$, the following limits hold

$$\lim_{v \rightarrow 0} F(v) = \frac{\lambda}{\sigma^2} - \ln\left(\frac{\lambda}{\sigma^2}\right) - \frac{1}{Q} - 1 \quad (\text{A.5})$$

$$\lim_{v \rightarrow (\pi/\lambda)} F(v) = \infty \quad (\text{A.6})$$

The first limit can be found by considering the following form for $F(v)$

$$F(v) = \frac{\lambda}{\sigma^2} - \ln\left(\frac{\lambda}{\sigma^2}\right) - \frac{1}{Q} - \frac{v\lambda}{\tan(v\lambda)} + \ln\left(\frac{v\lambda}{\sin(v\lambda)}\right) \quad (\text{A.7})$$

and noting that both

$$\lim_{x \rightarrow 0} \frac{x}{\sin(x)} = 1 \quad (\text{A.8})$$

and

$$\lim_{x \rightarrow 0} \frac{x}{\tan(x)} = 1 \quad (\text{A.9})$$

The second limit can be found by writing $F(v)$ in the form

$$F(v) = \left(\frac{\lambda}{\sigma^2} + \ln(\sigma^2) - \frac{1}{Q}\right) - \frac{v\lambda}{\tan(v\lambda)} + \ln(v) - \ln(\sin(v\lambda)) \quad (\text{A.10})$$

and noting that as $v \rightarrow \pi/\lambda$ from below the following limits hold

$$\frac{\lambda}{\sigma^2} + \ln(\sigma^2) - \frac{1}{Q} = \text{constant} \quad (\text{A.11})$$

$$-\frac{v\lambda}{\tan(v\lambda)} \rightarrow +\infty \quad (\text{A.12})$$

$$\ln(v) \rightarrow \ln(\pi/\lambda) \quad (\text{A.13})$$

$$-\ln(\sin(v\lambda)) \rightarrow +\infty \quad (\text{A.14})$$

The second part of the derivation requires us to show that $F(v)$ is increasing on the interval $v \in (0, \pi/\lambda)$. This can be shown by examining the derivative

$$F'(v) = \frac{1}{v} - \frac{x}{v} \left(\frac{2 \tan(x) - x \sec^2(x)}{\tan^2(x)} \right) \quad (\text{A.15})$$

where $x = v\lambda$, which can also be written as

$$F'(v) = \frac{1}{v} \left(\frac{\sin^2(x) - 2x \sin(x) \cos(x) + x^2}{\sin^2(x)} \right) \quad (\text{A.16})$$

This derivative can be shown to be positive as follows. The denominator is clearly positive on the interval $v \in (0, \pi/\lambda)$, which is equivalent to the interval $x \in (0, \pi)$. Examining the numerator

$$h(x) = \sin^2(x) - 2x \sin(x) \cos(x) + x^2 \quad (\text{A.17})$$

we see that this is also positive on the interval, since

$$h(0) = 0 \quad (\text{A.18})$$

and

$$h'(x) = 4x \sin^2(x) \quad (\text{A.19})$$

$$> 0 \quad (\text{A.20})$$

It follows that, since $F(v)$ is increasing to $+\infty$ on the interval $v \in (0, \pi/\lambda)$, a root of $F(v)$ exists on this interval, for a given Q and λ , when its lower limit on the interval is negative. From Equation (A.5), this corresponds to

$$\frac{\lambda}{\sigma^2} - \ln\left(\frac{\lambda}{\sigma^2}\right) < 1 + \frac{1}{Q} \quad (\text{A.21})$$

Now, as seen in Figure A.1,

$$\frac{\lambda}{\sigma^2} - \ln\left(\frac{\lambda}{\sigma^2}\right) \geq 1 \quad (\text{A.22})$$

with a minimum at 1 when $\lambda = \sigma^2$, and it crosses $1 + 1/Q > 1$ just once above $\lambda = \sigma^2$ and once below it. Outside of the open interval, bracketed by these crossovers, we have

$$\frac{\lambda}{\sigma^2} - \ln\left(\frac{\lambda}{\sigma^2}\right) \geq 1 + \frac{1}{Q} \quad (\text{A.23})$$

and thus $F(v)$ cannot have a root on $v \in (0, \pi/\lambda)$. It follows that these crossovers are the *minimum* and *maximum* eigenvalues for the exponentially weighted random matrix. Thus,




Figure A.1: $x - \ln(x)$

the theoretical maximum eigenvalue is the solution of

$$\frac{\lambda}{\sigma^2} - \ln\left(\frac{\lambda}{\sigma^2}\right) = 1 + \frac{1}{Q}, \quad \lambda > \sigma^2 \quad (\text{A.24})$$

We note also that Potters et al. (2005) have demonstrated an alternative derivation, with $\sigma = 1$, using “Blue” functions.

A.3 Summary

In this appendix we have derived an efficient method for estimating the maximum eigenvalue of an exponentially weighted random matrix. This alternative method offers a valuable saving of computation time compared to the original formulation.

Appendix B

List of Publications

1. Daly, J., Ruskin, H. J. and Crane, M. (2006). Exponentially Weighted Volatility, Random-Matrix-Theory Filters and Stability in Portfolio Optimisation, Oral Presentation to *Econophysics Colloquium 2006 and Third Bonzenfreies Colloquium, Tokyo, November 2006*. Publication in *J. Econ. Interact. Coord./Physica A*: cancelled due to budget restrictions. (Abstract bound in.)
2. Daly, J., Crane, M. and Ruskin, H. J. (2008). Random matrix theory filters in portfolio optimisation: A stability and risk assessment. *Physica A*, 387, pg. 4248. (Bound in.)
3. Daly, J., Crane, M. and Ruskin, H. J. (Submitted, 2009). Random matrix theory filters and currency portfolio optimisation. Forthcoming in *J. Phys. Conf. Ser.*, as part of the proceedings of APFA7: Applications of Physics in Financial Analysis 7, Tokyo, March 2009.
4. Daly, J., Ruskin, H. J. and Crane, M. Stability-based filtering of foreign exchange portfolios. In preparation, for submission to *Phys. Rev. E*, 2009.
5. Daly, J., Ruskin, H. J. and Crane, M. Can random matrix filters be used for trading in the foreign exchange market? - A comparison of foreign exchange and stock portfolio filtering. In preparation, for submission to *Dynamics of Socio-Economic Systems*

[illegible]

[illegible]

[illegible]

[illegible]

[illegible]

[illegible]

[illegible]

[illegible]

[illegible]

[illegible]

[illegible]

[illegible]

[illegible]

[illegible]

[illegible]

Appendix C

Matlab[®] source code disc