

A Disaster Waiting to Happen

Robert MacRae and Chris Watkins

Estimating the future risk of a portfolio is one of the basic tasks of financial analysis. Portfolio risk estimation is widely applied, for example in:

- Estimating VAR for management control of risk
- Designing portfolios to track an index benchmark
- Portfolio optimisation

The most basic approach is to use the observed, historical risk of a portfolio in the past as an estimate the portfolio's risk in the future. But using historic risk as a direct estimate of future risk is not valid in any of these applications, and leads to a systematic and potentially disastrous underestimate of risk.

We will show that this underestimate is due to selection bias. Explaining the mechanics of the problem allows us to suggest solutions -- and quantify what can happen if the problem is ignored. This article considers techniques that are applicable to all asset classes. However for simplicity we only consider portfolios of simple securities, not options.

An Introduction to Selection Bias

Selection bias can affect any type of statistical estimation, so let's discuss it in a non-technical environment, a golf tournament.

To estimate somebody's future average score on a course, we can have him play a few rounds: his average score is then an *unbiased* estimate of his likely scores on future rounds. Unbiased means that estimates made in this way are not systematically either too high or too low.

Suppose we organise a golf tournament in which 100 people play a round. Is the score of the tournament winner a reasonable estimate of his future scores? If the field consists of players of similar abilities, then the winner is likely to be one of the better players, but he will also have to be lucky on the day. The winner's score is likely to be above his personal average, for if he had not been lucky, he would not have won. This is selection bias. We have selected a player based on how good his score is, so his score is a biased estimate of his typical scores.

The amount of selection bias is affected by:

- **The number of competitors**
The more competitors, the luckier the winner will be. If the entire golfing population of the world were to play in a tournament, the professionals would, as a group, do well, but it would be unlikely that a professional would win. Having so many competitors would ensure that some lucky unknown won the tournament after playing the best round of his life.
- **The length of the competition**
If the competitors' scores are calculated from many rounds of play, rather than just one, each competitor's average score will be a more precise estimate of his ability, so that the bias will be less. For example, if we were to force all the world's golfers to play, say, thirty rounds of golf instead of just one, then it is likely that the winner of this tournament would be a top professional.
- **The distribution of ability among the competitors**
The narrower the range of players' true abilities, the more selection bias there will be. There are two extreme cases. If all competitors have exactly the same level of ability, then the difference of the winner's score from the average of all competitors is entirely selection bias. If, on the other hand, one competitor is sufficiently better than the others to be almost sure of winning, then he does not need to be lucky to win, and there will be no selection bias.

Selection Bias in Risk Estimation

We have shown that, when you select an individual on the basis of a score, it is not safe to use the score he achieved as an estimate of future scores. In risk estimation, exactly the same reasoning tells you that if you optimise a portfolio for minimum risk (equivalent to selecting it from among all the other possible portfolios), your estimate of its risk will be biased downwards. The portfolio will be riskier than you believe.

Lets examine this argument more closely. The analogy between golf and risk estimation (aside from the damage done by a really wild swing) is that each possible portfolio is a "golfer", and the portfolio's "tournament score" is its level of realised risk on the historic data. To be picturesque, we may imagine running a tournament by explicitly writing down a list of competing portfolios, and after each day of historic data, incrementing the score of each portfolio by its squared return so that we build up an estimate of its variance. The tournament ends when the historic data is used up, and the winning portfolio is chosen to be the one with the smallest score -- that is, the portfolio with the lowest historic volatility.

As in the golf tournament, the risk of the winning portfolio will be an *underestimate* of its future risk, because the winning portfolio must not only be low risk, but it also needs to be "lucky" in the sense that it has a lower than usual score in the particular historic data used. Just as in the golf tournament, the size of the estimation error will depend on the number of possible portfolios implicitly entered for the tournament, the amount of historic data used, and the spread of underlying risks among the portfolios.

The second part of this paper will discuss the magnitude of the possible errors, and the steps that can be taken to avoid them, but first it is worth pointing out just how

widespread this problem is. The argument above applies only to *optimised* portfolios, and most practitioners will be familiar with the problem in this context. Surely risk estimation on unoptimised portfolios will be immune?

Portfolio Risk Control

Unfortunately, unoptimised portfolios are very, very rare. Few fund managers have such faith in efficient markets that they pick stocks at random -- nor do we advocate this. Real portfolios tend to have been *selected* with some purpose in mind. Frequent examples are to "be conservative", to "avoid speculative stocks" or even just to "look sensible". All of these reasonable motivations involve some selection for low historic risk, and consequently the historic risk of the resulting portfolio is likely to be an under-estimate of its future risk.

Selection is present in almost all portfolio construction, and the threat of selection bias means that we should consider rather carefully how much selection is going on because the stronger the selection, the stronger the resulting bias is likely to be. In some cases, VAR models will be subjected to exceptionally strong selection by the traders who operate within them.

Company-Level VAR Control

A (modestly) cynical view of management risk control is that it exists to protect the interests of the firm against the interests of individual traders. Their interests differ when traders are remunerated with bonuses that are in effect call options on their book profits. Under these circumstances, traders have an incentive to increase both their expected return and their position risk, since the call is more valuable with higher volatility. The management risk control system exists to ration the amount of risk that each trader can take, so that they concentrate on the shared goal maximising expected returns.

Unfortunately, the control system is only as good as its risk estimates. Traders will seek to maximise their expected return within their VAR limit, by putting large amounts of capital into any trades for which estimated returns are high but VAR is low. Individual traders' incentive to maximise personal return subject to risk can subject the company's portfolio to intense optimisation that is not explicitly part of the organisation's goals, but is an implicit result of the trader reward structure.

Implicit optimisation acts rather like a mechanical portfolio optimiser, but there are several features of the process that make it a particularly dangerous form of portfolio optimisation.

- Since it is implicit rather than explicit it is hard to see in operation until a large loss is made. However, because the incentives are strong we can be rather confident that it is a common occurrence.
- Mechanical optimisers seek only to maximise a risk-averse utility function. Traders on the other hand have, through their call option, an incentive to be risk seeking so the implicit utility function may not be risk averse.

- If any trades carry too low a VAR tariff, all the company's traders will tend to take on exposure to them. This will lead to the management nightmare of correlated risks appearing across many nominally separate books. This has potentially dire implications for the whole risk estimation process.

Never underestimate the abilities of your traders, or the damage an able trader can inflict.

This provides a golden rule for managerial risk control systems:

Traders' position limits should never be stated purely as a limit on VAR estimated from historic risk.

In practice, organisations have lots of other controls as well. Simple constraints like restricting the geographical range, or range of security which can be held in a book make it much less likely that the same bets will appear in multiple books. The role of experience and common sense is even greater. However, it is important to point out that the current fashion for big, formal, all-embracing VAR systems would rapidly lead to disaster if it was allowed to replace, rather than support, these existing controls.

Even though these other controls act to limit the impact of the implicit optimisation, its inevitable presence means that there is always a risk that the company VAR estimate will be too low unless steps are taken to allow for this effect.

This line of reasoning provides strong support for the Basel committee's recommendation that any estimate of risk based on internal risk models must have a substantial safety margin added. Past validation or backtesting of an internal risk model cannot guarantee that the combination of selection bias and implicit optimisation will not invalidate it in the future.

The Basel Committee is right

Conclusion

So far we have shown that any practical estimate of future risk is likely to suffer from selection bias. The conclusion -- that once a model has been estimated, some allowance has to be made for worse performance out of sample -- should be familiar to every practitioner, but by making the reasoning explicit we can begin to quantify the severity of the problem and provide a solution.

The similarity between mechanical portfolio optimisation and the implicit optimisation carried out by traders suggests that, to be conservative, we should treat every risk estimate as cautiously as if the portfolio had been optimised for minimum risk.

For the worst case of selection bias, examine a fully optimised portfolio.

If a portfolio optimised using the risk-estimation system has out-of-sample volatility similar to the predicted volatility, selection bias will not be a problem. If out-of-

sample volatility is a substantial factor higher, then selection bias may inflate experienced volatility by up to this factor. This approach will be developed in the next two articles.

To Follow:

We will use the approach of examining optimised portfolios to:

- 1) Show how selection bias depends on the three critical factors (amount of data, number of alternatives, distribution of risks).
- 2) Construct three examples (two synthetic, one real-life) that demonstrate the real-life severity of the problem.
- 3) Provide a test that can be applied to the covariance matrix associated with any VAR model to show how vulnerable it is to selection bias.
- 4) Provide a simple and widely applicable method for limiting the potential damage.
- 5) Illustrate this technique with the three examples introduced earlier.

A Disaster Waiting to Happen: Modelling and Analysis

In the first article we introduced the concept of selection bias and showed that it is liable to corrupt any practical estimate of future risk. The concept of implicit optimisation was introduced to suggest that most real-life portfolios have been optimised to some unknown degree, and that this poses a significant risk to all enterprise-wide VAR models.

In this article we examine in a quantitative way portfolios that have been optimised for minimum risk, which are the worst case. Rather than relying on extensive maths we will construct three examples that can be duplicated in any statistical package or in a spreadsheet so that the reader can test these ideas.

We will start with an artificial example, that of constructing a hedged long-short portfolio of uncorrelated securities. This permits us to make some general statements about how bias varies with number of stocks and length of data.

This example is chosen purely to make the model simple; it is very far from being a worst case. The same problems lurk, hidden, in all portfolios of all securities and deviations from normality, missing data and autocorrelation all tend to exacerbate the problem. The second example generalises the first by including known but different volatilities, and the third example uses real data on currencies to show the problem in a simple but realistic context.

In the final part of the article we will deal with methods that can be used to diagnose and improve the situation with the emphasis on simple, practical steps rather than theoretical derivation.

Modelling: computational experiments

A) Independent Assets

We have stated that the severity of selection bias increases with the number of available choices (the degree of selection) and decreases with the amount of data. We want to examine the strength of these effects. Is the problem serious with practical amounts of data?


We generated artificial random returns data for various numbers of securities with the same volatility and uncorrelated with each other. Such data is particularly easy to generate, and we recommend the sceptical reader to try the following experiment.

One advantage of artificial data is that we can generate as much as we need, and can examine many different sizes of problem. We chose to examine the range of 5 to 50 assets, with a number of observations ranging from 1 to 50 times the number of assets. We can also calculate the theoretical volatility of any portfolio, and any deviation from this will be due to the finite data available for estimation. To permit

comparison of portfolios with different numbers of assets, we normalise the sum-of-squares position sizes to give all portfolios equal long-term volatility of 5%.

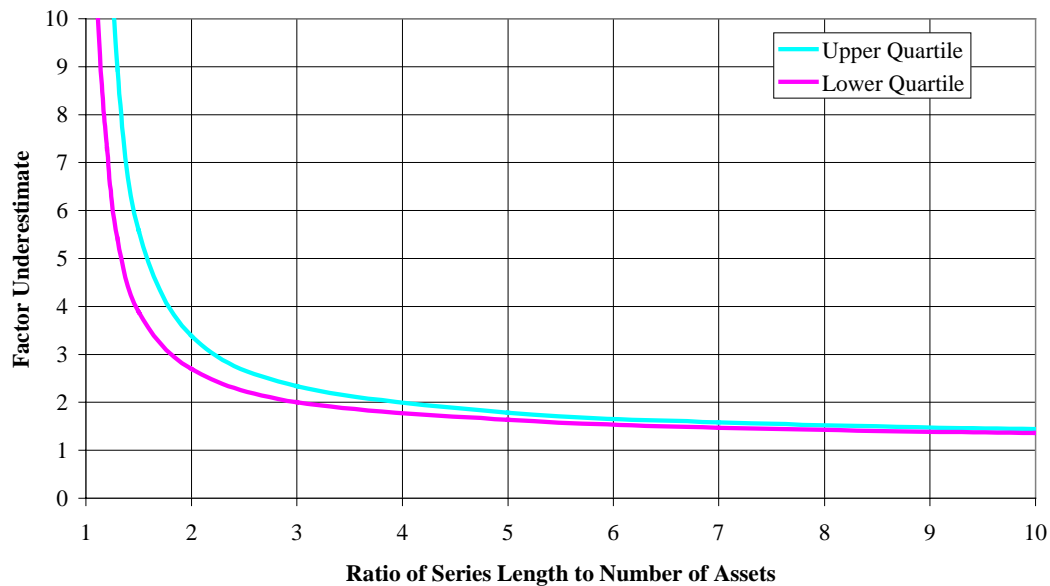
For each set of data we construct the covariance matrix and invert this to find the portfolios with maximum and minimum volatilities. Figure 1 is generated by replicating each case many times and plotting the average values for both maximum and minimum risk portfolios.

Fig 1: Volatility Estimates as a Function of Data to Assets Ratio



In practice we are only concerned with low-risk portfolios, for which the error is almost independent of the number of assets involved. This permits us to construct Figure 2, which shows quartiles of the ratio of out-of-sample to in-sample volatility for the minimum risk portfolio.

Fig 2: Volatility Underestimate for Optimised Portfolios due to Finite Data



There are three key conclusions from these graphs:

1) Low risk is proportionally worse estimated than high risk

Perhaps the most ironic result of the demonstration is that the bias on estimates of lowest risk are much worse than on estimates of high risk. It is unfortunate that it is precisely these hard-to-estimate low-risk portfolios that are of most practical interest.

2) The amount of historic data needed is proportional to number of assets.

If you have ten times as many data points as you have assets, selection bias leads to an average volatility underestimate of about 30%, which is not unreasonable in comparison to the other uncertainties of the risk-control process. However, if you use only twice as many, selection bias will probably lead to out-of-sample volatility being 2.5 to 3.5 times as large as you estimate. This result has been found to hold for much more general artificial cases including those with a shared "Market" factor, and APT structure, and we believe it to be a good general guide. It tends, however, to be *optimistic* with long-tailed and heteroskedastic distributions, since relatively few of the observations make any significant contribution to the covariance and the effective number of data points is often far lower than the total number.

3) Sufficient historic data is often unavailable in practice.

At first sight, a requirement of 10 data points per asset does not appear too demanding since daily data is widely available, but there are pitfalls in using high-frequency data. These problems are really outside the scope of this article, but it is worth pointing out that daily data on less-liquid securities often exhibit substantial non-trading, bid-ask bounce and other short-term effects that lead to substantial errors if daily volatilities are extrapolated to longer time-scales using the familiar \sqrt{t} assumption.


Introducing another rule of thumb, if you are interested in forecasting on a certain time-scale (for example 1 year) there are limited benefits to using data that is spaced more closely than 1/10 of this time-scale (since shorter term effects are likely to dominate and closer points cease to be independent), or greater than 10x this time-scale (since non-stationarity becomes a problem and older points cease to be relevant). If we were to take these guides as hard limits then we would be limited to 100 independent and relevant data points per security -- so we could safely estimate risk for optimised portfolios containing at most 10 securities!

This conclusion is over-pessimistic, but it should serve as a warning that any practical optimisation involving hundreds of securities cannot be assumed to be immune from selection bias simply because several years of daily data are available.

B) Known Unequal Volatilities

The simple example above cannot demonstrate the effect of the distribution of volatilities since by construction they are all equal. Figure 3 was created by constructing a covariance matrix for 10 assets that shared a substantial market factor with a 20% volatility, and residual volatilities evenly spaced between 6 and 10%. These theoretical volatilities are shown as the asymptotes to the right of the graph. We sampled finite data from the covariance matrix to estimate these values for various numbers of data points, giving rise to the curves to the left.

Fig 3: Estimates of Unequal Volatilities



Note how well-estimated the market factor is; there are no other factors of similar volatility so there is little selection bias -- in terms of our golf metaphor, it is the sole professional in an amateur competition. By contrast, the 10% volatility is significantly overestimated and the 6% volatility is underestimated, though for 100 or more data points these effects become minor.

C) Real-World Example

With real-world data we do not have the luxury of access to infinite amounts of data, so we cannot create charts based on multiple replications. However we can show the practical effects by looking at in- and out-of-sample performance. Almost any group of financial assets could be used for this illustration, but we will use 5 years of monthly dollar returns for 50 currencies so that we have 1.2 historical points per asset. Such a covariance matrix would be appropriate for forecasting long-term foreign exchange risks, and illustrates the practical problems of getting enough data. Though for most of these currencies we would be comfortable with a weekly sampling rate, the problem of relevance is particularly severe given the occasional dramatic shifts in exchange-rate regimes.

Currency	Portfolio Weight	SD
FRENCH FRANCS	.525	12.2%
NORWEGIAN KRONER	.380	11.8%
MALTESE POUNDS	-.347	10.7%
SINGAPORE DOLLARS	.278	4.1%
SWISS FRANCS	-.181	12.9%
GREEK DRACHMAS	-.179	11.3%
COLUMBIAN PESOS	.176	12.7%
KOREA (SOUTH) WON	.176	4.1%
PAKISTANI RUPEES	-.175	7.7%
SWEDISH KRONER	-.154	13.4%
BELG/LUX.CONVERTIBLE	-.144	12.5%
GERMAN MARKS	-.135	12.3%
IRISH PUNTS	.135	12.4%
SPANISH PESETAS	.127	13.9%
ISRAELI SHEKELS	-.120	9.1%

This table shows the 15 largest positions in the minimum-risk portfolio -- that is the portfolio with lowest dollar volatility. Only relative weights matter, so for convenience we have set the sum of squares position size to 1.0. For reference the volatility of the currency over the in-sample period is also shown. The portfolio is visibly "Hedged" in some sense, and several obvious sources of risk are at low levels:

- Net US Dollar exposure is only 0.146,
- Net European exposure is only -0.022,
- The five countries with volatility greater than 15% are amongst the smallest 12 position sizes.

However the enormous positions in French Franc and Norwegian Kroner do not look low risk, even hedged with a basket of other European currencies. The concentration of these positions makes the in-sample estimate of volatility, 0.26%, seem ridiculously low.

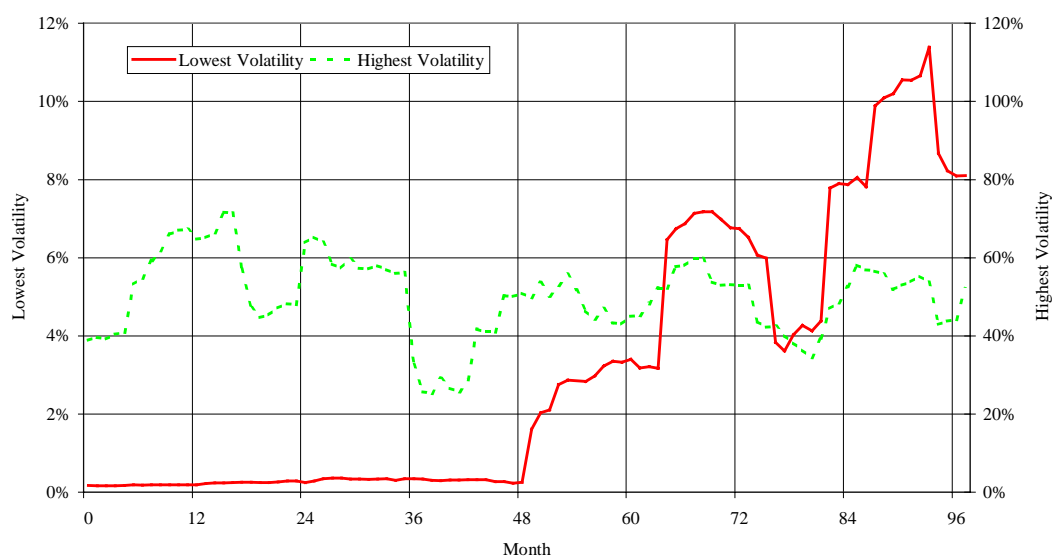
Volatility for Lowest and Highest Risk Portfolios

	Lowest	Highest
In-sample	0.26%	51%
Out-of-sample	7.43%	52%

In fact 0.26% is a disastrous 25x underestimate out-of-sample. The chart shows the 12-month lagged volatility for the optimised portfolio with months 0-47 in-sample and 60-97 fully out-of-sample. The unrepresentative nature of the in-sample period is very clear.

Recall what this means in practice. If you estimate the risk of your current positions using this historic matrix and your positions have been selected for low risk, then you will underestimate the risk of the position by an unknown factor of up to 25. This should worry you even if you believe that the degree of selection has been modest.

Fig 4: In- and Out-of-Sample Volatility



For comparison, the chart also shows the volatility of the most risky portfolio on the right-hand axis. Though it is always much riskier than the hedged portfolio, the in- and out-of-sample periods have similar volatilities. Selection bias is a minor problem for the most risky portfolio, with no significant difference between in- and out-of-sample volatility.

This example is small enough to replicate in a spreadsheet. Once data has been input, spreadsheet functions can be used to construct the covariance matrix and optimise a portfolio for minimum risk. The detail results will of course vary depending on your choice of assets and time periods, but the general conclusions hold for almost any similar group of assets.

Conclusion

These examples show that selection bias is a significant problem in realistic optimised portfolios. To avoid it when there are hundreds of assets in the optimisation, several thousands of independent data points are required. High-frequency data must be used with caution unless all assets are highly liquid, and old data is suspect due to changes of regime.

It was argued in the first section of the article that most practical portfolios have been subjected to some implicit optimisation. Since it is hard to estimate how extensive optimisation has been, it is conservative to look at the worst case, that of complete unconstrained optimisation. These examples suggest that selection bias will usually be a significant threat in practice. The practical example in which the volatility of an optimised portfolio is underestimated by twenty-five times is severe, but not unreasonable.

Fortunately, as well as indicating the severity of the problem the analysis of optimised portfolios also suggests steps to improve the situation in practice. This will be the subject of the final section of this article.

A Disaster Waiting to Happen: Prevention!

In the first part of this article we introduced the concept of selection bias and showed that it is liable to corrupt any practical estimate of future risk. This was demonstrated for optimised portfolios, and the concept of implicit optimisation was introduced to show that most real-life portfolios are affected and that this poses a significant risk to all enterprise-wide VAR models.

In the second part we examined both artificial and real-life examples to show that the problem of selection bias can be serious in practice and to show how it varies with the number of assets, length of data and distribution of volatilities.

In this, the final part of the article we will use the results of parts 1 and 2 to suggest methods that can be used to diagnose and improve the situation. Our approach is to suggest a range of simple, practical steps rather than to give a unified theoretical treatment. This is because the theory would have to deal with abstracted and idealised cases, with the attendant risk that when it came to be applied the theoretical assumptions would not hold. Risk control is a discipline that demands conservatism and robustness above all other considerations so we have a strong preference for simple techniques.

Diagnosis: Examine the Spectral Volatilities

Somewhere inside every risk-estimation process lurks a covariance matrix. It may be estimated directly from historical data. It may be estimated via an exogenous factor model as in one of the BARRA commercial packages, or via an APT approach which contains a limited number of unspecified factors. There are limitless variations on the theme of estimation to handle financial data features such as illiquidity, non-normality, heteroskedasticity, shocks and non-stationarity. The test of a good estimation process is its ability to handle these features.

Rather than attempt to provide a complete estimation package, we are going to describe how to check the extent to which the covariance matrix constructed by your existing package is prone to selection bias. It is far simpler to describe the test than to construct the theoretical derivation, so we provide the method first:

Extract the minimum spectral value of the covariance matrix. This is the variance of the minimum risk portfolio associated with the covariance matrix. If it is lower than you consider plausible, you have a problem.

Argument

Technically, spectral value decomposition is a matrix decomposition that splits an N-by-N square matrix into two orthonormal N-by-N matrices and a diagonal weight matrix. Most textbooks on linear algebra include some material on it, though

practitioners may prefer the more practical approach of Numerical Recipes. Practically, it is the decomposition that powers Principal Component Analysis. It finds the largest contribution to variance, which is the first principal component, and then proceeds to find successive components in decreasing order of importance until all the variance of the matrix has been explained.

In the financial case, decomposing the covariance matrix in this way constructs N orthogonal (or uncorrelated) portfolios of known variance. The sum of squares of the asset sizes in each portfolio sum to 1.0. Because of the orthogonality it is easy to decompose any portfolio into a linear combination of the portfolios (which we term "Spectral Portfolios" from their association with the decomposition), and its variance will be the same linear combination of the "Spectral Variances".

If you wish to understand the covariance structure, it is much easier to work with these spectral portfolios than the matrix itself. For example, if you plot the highest-variance portfolio you will usually see that it is exposed to an obvious market factor, while the low-variance portfolios are all hedged.

In fact, extracting the minimum spectral variance is exactly the same as optimising a portfolio for minimum variance subject to the constraint that the sum of squares of the positions sum to 1.0, and we have shown that examination of this optimised portfolio permits us to estimate the potential severity of selection bias. The formalism of Spectral Value Decomposition is, however, going to be essential for constructing a cure.

Returning to our prescription above, if you extract the minimum spectral volatility, this is the volatility of the minimum risk portfolio. Is it reasonable? To answer this question it may be helpful to look at the portfolio associated with it and to ask what volatility you would expect out-of-sample. This may sound like a rather tricky question, but in practice the answer does not have to be very accurate. In the currency example we examined earlier, the volatility of 0.26% was clearly wrong. It is often the case that the minimum spectral volatility is effectively zero for portfolios no one would regard as risk free. In this case the problem of selection bias is serious whether the true volatility is 1% or 10%.

The method usually provides an appropriate diagnosis. It is possible to construct examples that require some further analysis -- for example by including assets which span a very wide range of volatilities, or which include arbitrage relationships -- but the principle of decomposing and then asking whether you believe what you see is applicable to all cases. That just leaves the problem of what to do about it.

Cure: Fix the Spectral Volatilities

In fact, this method of diagnosis suggests the cure. The key step is to remove the aspects of the covariance matrix that we consider implausible. Again, the method is simpler to describe than to justify, so the treatment comes first.

Extract the spectral values of the covariance matrix. If any are lower than you consider plausible, increase them to a plausible level and then re-compose the covariance matrix.

Argument

Proof of optimality is tedious and requires assumptions, such as normal distributions, which are not met in the most important cases. We prefer various forms of hand waving to suggest that this approach is reasonable. For a start, it is obvious that a covariance matrix edited in this way will pass the diagnostic test given above. We also argue that this modification will usually be inoffensive. If you compare risk estimates using the edited matrix to those using the raw matrix, you will find very minor differences for most portfolios; all portfolios will appear to have the same risk as before, or more, but unless the "plausible value" you entered is almost as large as the largest spectral values, the change in *most* risk estimates will be minor due to the sum-of-squares addition of volatility. If the value you have entered is small compared to other sources of variance then its overall contribution will be minor.

The few portfolios significantly affected are those that, before, had implausibly low estimated volatilities. The procedure of editing the spectral variances, which we refer to as "Plausibility Editing" guarantees that no spectral portfolios have implausible volatilities. It can easily be shown that all their linear combinations must also be plausible. Since estimates are now all guaranteed to be larger than the smallest volatility you provided, there will be little to fear from optimisation, implicit or otherwise.

Finally, consider some other, competing technique that achieves a similar matrix conditioning by a different method. If, after it has been applied, our decomposition approach yields some spectral volatilities that you consider implausible, you should apply our approach as well to remove them. Other methods may be superior in other senses (such as building valuable priors into the matrix), but we are not aware of any that have the robustness of our approach. Once you have done everything else, edit the plausibility!

This is a simple yet highly effective cure for selection bias. The major advantage of our approach is that practitioners can use their market knowledge to set this number without the need for any statistics. "What is the lowest volatility estimate you would trust?" is a particularly intuitive question to ask, and the spectral value decomposition lets us build the answer directly into the covariance matrix.


Plausibility Editing in Practice

Applying this technique to the examples mentioned earlier illustrates the general nature of the impact. Again, we would urge the sceptical reader to try out this approach on his own data since the essence of this technique is its breadth of application.

A) Identical Independent Assets

In this artificial case we know that the minimum spectral volatility should be 5%. However, in practice we would only have an estimate of the true level. Figure 5 illustrates the effect of imposing a floor of 2.5% on the spectral volatilities. As expected, no portfolios now appear to have volatilities below 2.5%, so we have removed the possibility of a dangerous volatility underestimate. There is no significant inflation of estimated volatilities that are above 3%.


Fig 5: Impact of 2.5% Volatility Floor on Equal Volatilities



B) Known Spectral Values

In the last section we imposed a floor that was half the theoretical value. In practical cases we never know exact values, but the method is reassuringly robust to imprecision in the estimate. For the second example we set the plausible floor equal to 5%, which is just below the lowest of the theoretical spectral volatilities. There is little or no impact on volatility estimates above 6%, but no possibility of an estimate below 5%.

Fig 6: Impact of 5% Volatility Floor on Unequal Volatilities



C) Real World Example

Here we have no theoretical value on which to base the floor, so we have to use other methods. Resampling is our preferred statistical tool for this but because the only requirement is plausibility we do not really need this sophistication. An argument that puts us on the right general scale is:

"Many currencies have a volatility around 10%, but the presence of correlated trading blocs means we have access to only 4 such assets which are truly independent. We do not wish to assume that the blocs are stable and so are unwilling to take hedged positions within a bloc, but we can diversify four ways which suggest a volatility of $10\% / \sqrt{4} = 5\%$."

To the extent that you trust the stability of currency blocs you might prefer a lower floor, perhaps 2% or even 1%, allowing a greater role for hedging and consequently accepting less security. However, the precise level of the floor is not critical.

Imposing the floor of 5% reduces the potential underestimate from 25x to 1.5x, which reduces selection bias from threat to nuisance, and imposing a floor of 2% would reduce selection bias to 4x. Recall what this means in practice. If you base your VAR system on the edited matrix then to whatever optimisation process it is exposed, or however smart a bunch of traders abuse it, you expect out-of-sample volatility to be at worst 1.5x the in-sample estimate rather than potentially 25x if you had used the original matrix.

This example demonstrates that the method does not depend critically on the choice of plausible floor. So long as the value is just that -- plausible -- there will be a dramatic improvement in performance in many real situations. If the floor is set too high, the volatilities of well-hedged portfolios will be overestimated and the company may take on less exposure than it could. If set too low, the power of hedging will be overestimated and selection bias will be reduced but may remain a problem. There is a broad range of values for which the gains are substantial.

Practical Considerations

This is a powerful technique. Such techniques carry a particular risk that they are seen as a panacea and lead to insufficient attention being applied to other areas. We would note that in addition to using this technique you should still:

- Prepare your matrix with care, incorporating all the prior information you possess.
- Simulate the end-to-end results and check against reality.
- Ensure that non-VAR position constraints are in place, and sufficiently binding to provide an independent source of security.
- Watch your traders like a hawk.

We would also like to point out that this method has an interesting and beneficial interaction with position constraints. Many practitioners have observed that optimisation of positions subject to constraints often leads to over-trading, and minor changes to the expected returns or the covariance matrix leads to substantial trading, even in the presence of realistic transaction costs. This behaviour can be shown to be due to the presence of unrealistically small spectral volatilities, and disappears if these have been increased to plausible values. The end result is that constrained optimisation combined with edited spectral volatilities does not lead to over-trading.

Conclusion

Risk control is a complex field that has an endless capacity to hide problems, for example in data collection, data quality and modelling. The only safe way to combat this complexity is to make assumptions and models as explicit as possible so that they can be examined. We have demonstrated this approach on one of the key components of the risk-control problem, the covariance matrix of asset returns. Plausibility Editing ensures that we are not building implausible assumptions into our estimates of risk.

This article pursues two key threads. The first is that selection bias is an insidious problem that affects most practical risk estimation problems as a result of portfolio construction methods and implicit optimisation. This indicates that any matrix that would be unsafe for use in an optimisation will also be unsafe, to some unknown degree, when used for risk estimation.

The second thread provides a simple technique that can be used to make any covariance matrix safe for optimisation purposes. This is useful in its own right, but combined with the first thread it takes on a very broad applicability. Plausibility Editing should be applied to most covariance matrices used anywhere in the risk estimation process.