# GRESHAM COLLEGE

EST. 1597

## The folly of Value-at-Risk
## Professor Avinash Persaud
### 2 December 2002

## Introduction

Somewhere in the last century finance theory made a wrong turn and ended up in an intellectual cul-de-sac. Today, modern finance elegantly describes a world that does not exist. It assumes that investors behave independently when we know that strategic behaviour is rife, and herds are common. It assumes that individual investors' appetite for risk is very different and is fixed over time when their behaviour appears similar and their appetite for risk varies over time. It assumes that market participants have an independent notion of value, when valuation is highly uncertain, often self-fulfilling, and sometimes indeterminate.

Of course, economists have made a profession out of faulty assumptions. We argue, or perhaps hope, that precision in assumptions matters only if it changes the conclusion. It is therefore disturbing that, if we were to replace these faulty assumptions in finance with what we know, we reach some perverse conclusions, namely that the modern practice of finance does not reduce risks and may even contribute to market instability. It is certainly the case that despite developments that have been kind to the precepts of modern finance - i.e. regulation, more information and more adherence to quantitative financial analysis - the financial world is not markedly safer than before.


### Chart 1: Currency crises through time

Across a broad sweep of time, the number of currency crises appears to be rising. Eichengreen and Bordo have calculated the number of occasions that currencies have come under intense pressure in developed and developing countries across the last century - see Chart 1. Using more high-frequency, up-to-date data, market volatility across the largest equity markets also appears to be on the rise.


### Chart 2: Number of days per quarter that major markets have moved by more that 2-standard deviations in 3 major equity markets

This chart shows the number of days in a quarter that the US, Japanese and UK markets have moved by more than two standard deviations.

But is volatility a problem? I am not of the mind that all volatility is bad. Indeed, it is a prerequisite of a market economy, but policymakers should be concerned that it is of manageable proportions and those who can do so are managing that it. One measure of whether volatility is too high is whether it is acting as an obstacle to investment and borrowing. There is little evidence of this for the largest companies and richest countries. But there is evidence that volatile markets are not good at providing a flow of funds to emerging companies or countries - and while the periphery does not have a powerful lobby, it is important to remember that global growth will ultimately come from these quarters. These borrowers have to depend more on the fickleness of bank finance wherever and whenever it is available. Emerging countries represent well over two-thirds of the world's population, around one third of the world's GDP, but just 6% of the world's market capitalisation.

Some of you have heard me talk before about the role of modern risk-management systems in creating risks, but today I would like to do two new things. First, I want to show some new empirical evidence for my heretical claims, and then I want to use this evidence to identify what is wrong with modern financial risk management and to develop a new framework for thinking about financial risk.

Modern teaching on finance tells investors to do three things. The first is to define financial instruments, like stocks, bonds and currencies, by their average returns, volatility and covariance. Volatility of individual instruments and their covariance are key to the risks of a portfolio. Covariance is the degree of co-movement of two instruments when they vary from their average. A high positive number means that the two instruments move away from their average in the same direction at the same time. When it comes to picking your investment portfolio, high covariance is a bad thing; it is the equivalent of putting all your eggs in the same basket before you go on a long, bumpy journey.

The second thing investors are supposed to do is to find the optimal portfolio by picking instruments that as a group have high average returns, low volatility and low covariance. The third thing they are supposed to do is to manage their selected portfolio within a risk budget. If volatilities and covariances rise, risks rise. If the risk budget is breached, investors are required to reduce risk by selling volatile or highly correlated instruments. There is an aura of science about all this and economics in general, and finance in particular, have always been desperate to be counted among the harder sciences - at times too desperate. Adam Smith, the father of modern economics, was born four years before the death of Sir Isaac Newton and a Newtonian approach to the subject can still be observed today.

Let us now look at what happens in practice. Because of the need to compare on a like-for-like basis, fund managers generally group the portfolios they manage into three categories: low, medium or high risk. Most company pension schemes ask their employees to tick which bucket they want to be in. Of course, as John Nugee and his colleague Alan Brown of SSGA never tire of pointing out, risk is being poorly defined here. It should relate to the likelihood that your investments will not meet your financial needs when you need them to do so, but it most often refers to the volatility of returns relative to the performance of some broad stock index. A 59-year-old in the UK with her life savings in a so-called low-risk index-tracker fund would have found that the value of the pension she could draw for the rest of her life had lost 40% within the space of the first six months of this year. This does not seem to be low risk for her.

Leaving this aside, let us imagine that we are all fund managers in the medium-risk category, typically defined. Speak to any executive at Reuters, or former executives at Telerate or Knight Ridder, and they will tell you that financial data have become a commodity. We all have five or more years of data on the returns, volatilities and covariances for a similar universe of financial instruments, and so, given a similar risk budget and a similar universe of instruments, our optimised portfolio will be similar, if not the same. This will be fine, initially. Because new money is being put to work overtime, the assets in our optimal portfolio would probably be outperforming, adding to our confidence that we had made the right bets.

The problem is that the process whereby investors identify and then select the same optimal investment portfolio changes its nature. The assets in the portfolio will no longer be high-return, low-volatility and low-correlation assets, but the precise opposite. Because we have all bought them in search of the high returns of the past, they will become overvalued assets, incapable of outperforming others in the long run and vulnerable to bad news.

### Chart 3: VaR = Viciously Augmented Reversals

When bad news strikes one of the assets that we have overweighed in our portfolio and it falls in a volatile manner, we will quickly overstep our risk budget. As a result we will have to reduce risk by selling this asset and other assets in the portfolio. These are assets that we chose because they were not correlated, but they will now become correlated as we sell them together. When price declines bring out more sellers and not bargain hunters, volatility rises further. The rise in volatility and covariance causes a further breach of the risk budget and more selling. The portfolio selected for its history of high returns, low volatility and low correlation instruments has convulsed into a portfolio of negative returns, high volatility and high correlation. Instead of feeling smugly confident of our scientific method, we now feel a bit daft. It is also a delicious paradox: the act of observing an area of financial safety made it risky.

This is not such a delicious paradox for those financial academics or risk consultants or risk managers who feel threatened rather than energised by the intellectual challenge of this observation. Recently a colleague passed on a paper written by an academic who is also the author of a popular handbook on how to apply

2

the standard mean-variance and value-at-risk approaches I have just described. While he forgot to add me to his mailing list, the paper singles out my work for criticism and argues that the world is a safer place thanks to the use of the techniques he sells. For his sake let us go beyond the hypothetical and look at some real examples.

### Chart 4: Five year covariance matrix, FX Markets, 1995/96

This is the covariance matrix of currencies in the summer of 1995 with the markets grouped geographically, running from left to right as European, Latin American and Asian markets. I have also added another line at the bottom that shows the five-year average returns of these currencies in dollar terms. There are a lot of small numbers here, so I have coloured purple the characteristics investors want, which is high returns and low covariance and I have coloured red the characteristics they are trying to avoid, low returns and high covariance. In the summer of 1995-96, in the shadow on the 92-93 EMS crisis, European currencies did not appear the place to be: returns were low and covariance high. The place to be was Asia. Follow the purple patch. Asian markets not only offered some high returns, but they were uncorrelated with European and most American markets. So investors piled in. Not unwittingly, not unthinkingly, but in the intellectual footprints of finance's finest.

### Chart 5: Chart 4 with a colour key, purple for high returns and low covariance and red for low returns and high covariance, followed by 6 month covariance and return June 1998

We all know what happened. Here are the matrices of covariances and returns in the first six months of June 1998. Notice the metamorphosis. The purple patch has turned red. Asian currencies are now delivering not just underperformance, but the most negative returns, highest volatility and correlation. Was this just an unfortunate coincidence?

### Chart 6: Covariance matrix of global equity sectors, 1992-1997

Back in the mid-1990s, investors began to notice that global equity sectors had low covariances with one another and so offered better diversification than countries. So investors began to focus on sector investing and began to ignore country factors. Given that sectors had such low covariance (notice the swaths of purple in Chart 7) they could concentrate on the sectors with the highest returns, such as technology. But what they thought was a portfolio of assets with high returns and low covariances in 2000 turned into a very different portfolio, one that was not only going down, as most things were in 2000, but was going down in a much more correlated fashion than their VaR models predicted. The purple patch was "reddened". Sectors that were traditionally and economically separate, like financials and technology, or industrials and consumer discretionary, were acting as one.

### Chart 7: Covariance matrix of global equity sectors, 2000

As investors reacted to the observation of greater diversification by chasing after sector bets and avoiding country bets, sectors became less diversified and countries more so. Year-to-date the difference between the best- and the worst-performing major markets is as high as 20%. In today's investing, countries are back as a factor.

### Chart 8: Market diversification

Like country factors, emerging markets have gone from being discarded as a separate asset class after the 1997-98 Asian Financial Crisis to being today's hot sector. Strongly negative and highly correlated returns in both 1997 and 1998 dominated the returns-covariance matrices in the years following and kept investors out of emerging markets. Many emerging market managers lost funds and some closed down. Consequently, while contagion from the Argentine default, the largest default in history, has been deep and troublesome in a few places, notably Argentina's neighbours such as Brazil and Uruguay, the absence of a crowd of investors bailing out has meant its effects have not been widespread - certainly not as widespread as during the Asian financial crisis which engulfed over sixteen markets by the time of the Brazilian devaluation in January 1999. Relative safety in emerging markets has been born out of past risk. Today,

emerging markets as an asset class are outperforming with low volatility and covariance.

### Chart 9: Emerging market returns

I could continue like this all night, but I hope it is already clear that this paradox is robust across markets, sectors, asset classes and time. Let us pause to summarise exactly what we have seen and highlight the main principles at work so that we will have a better view of the solutions.

I am not saying that you could not find good fundamental reasons for investors' initially turning to Asia in the mid-1990s, or to the global technology sector in the late 1990s, or for their turning against these asset classes later. I am not just saying that correlations and volatilities are unstable, but that they are unstable in a very particular way. I have presented a profound paradox: the wide spread observation that a collection of assets are safe makes them riskier, and the observation that a collection of assets are risky makes them safer. This has interesting investment implications, and implications that are more than interesting for the current practice of risk management.

This paradox arises because of three aspects of investor behaviour. First, investors know that past returns, volatility and correlations are not fixed, but they view them as a given - as independent to whatever they do - as exogenous. Second, and partly as a result of collapsing computer costs, investors use similar information, similar technologies and similar preferences and so end up hunting a similar group of markets. Third, concentrated purchases or sales of markets changes their statistical properties - their likely return, volatility and covariance. Systemically, investors' reaction to data about financial instruments that they view as exogenous makes those instruments endogenous.

### Chart 10: The distribution of Asian covariances, 1992-1997

Understanding the process that is driving our paradox helps to explain why the proposed solutions to risk management mistakes have not worked and cannot work. There have been two principal proposals, of which the first is to have more sophisticated statistical models. In chart 13, I show the distribution of five-year covariances between pairs of seven Asian markets between 1992 and 1997. It is a narrow range of low numbers. Today, sophisticated risk managers assume that, while the future will be anchored by the past, in a crisis,this distribution may get fatter or more skewed.

Every time they fatten the distribution further we see the harried risk manager in the glare of a crisis complaining that his calculations were thwarted by a once-in-a-thousand-year event. These once-in-a-thousand-year events - like the dramatic drop in the Korean won in December 1997, the Russian default in August 1998, the LTCM liquidity black hole in October 1998, the Turkish and Argentine financial crises last year - seem to be occurring a bit more frequently than once in a thousand years. This is because the model is wrong: the future is not anchored by the past, but is a reaction to it. The more investors see and react to the old distribution of returns and correlations, the further away the new distribution will be from the old. We are not witnessing freakish once-in-a-thousand-year events every few years but the natural by-product of a different world than the one risk managers are modelling.

### Chart 11: The distribution of Asian covariances in 1998

The other main solution to the failure of modern risk management has been the use of stress tests. This is the assumption that the new distribution of returns, volatilities and correlations would be a repeat of what happened during some past crisis. It is a step forward not to view the future as an average of the past, but this will also be off the mark given the underlying process we have observed. If the markets you currently own were at the centre of the last crisis, the history of negative returns and high covariances would have kept other investors away and made the new distribution safer than the past crisis would suggest. The stress test will significantly overestimate your risks. More worryingly, if the markets you are in appeared safe during the last crisis, which is perhaps why you and everyone else are there, this concentration of positions will make the future distribution far riskier than during the last crisis. Repeating the past crisis on your portfolio is a distraction. What you need to focus on is how a reaction to the past will shape the future.

Interestingly, the current practice of risk management falsely assumes that there is no reaction by investors to the data they observe. This falsehood is revealed after each crisis. It is therefore time-inconsistent, and time-inconsistent practices will ultimately fall by the wayside.

So what can you do? I will describe how you should treat estimates of covariances but returns and volatilities can be treated similarly. Consider that covariances between assets in a crisis are made up of two elements: the long-run structural covariance between the two assets and a cyclical component. The structural element can be estimated by looking at the covariances of economic or corporate data such as current-account positions between countries or earnings between companies. This covariance of the fundamentals represents a minimum for the overall degree of covariance of market prices.

Next add the cyclical component. What you are trying to do with an estimate of the cyclical component is capture the impact of strategic behaviour and market positioning on future covariance. Ideally you would have data on market concentration, and the cyclical component of covariance would be zero when market concentration is low or so high as to make the overall covariance equal to one if market concentration was high. Of course, outside custodial databases it is hard to find good measures of market concentration. However, one of the lessons from the examples I have shown you is that you could estimate the cyclical component of covariances using past covariances, but in precisely the opposite way to how it is done today. Instead of assuming that the future is some positive function of the past, assume it is an inverse reaction to the past, so that low covariances in the past mean high covariances today and vice versa. This approach would be time consistent.

What I have been talking about is a failure of modern risk-management practice, but at its roots, and I do not say this lightly, it is a failure of modern finance. The science of finance takes from physics and engineering the one assumption that does not work in finance - that the players and instruments, or investors and markets, are independent of each other. In a reality of plentiful information but widespread uncertainty, strategic behaviour is rife: what you do changes what I do which changes what you do. This is mathematically inconvenient, and it draws us further into a world of indeterminacy and instability, sometimes called reality. But as John Nash showed us earlier, there are ways of addressing strategic behaviour. The solution I proposed above to the problem of estimating covariances is perhaps not elegant, but the more people start thinking about finance in terms of strategic behaviour, the more elegant solutions will surface.


**Conclusion**

I have presented a simple and disturbing paradox: the observation that a collection of assets is safe makes them riskier, and the observation that a collection of assets is risky makes them safer. The one time you need a good risk-management system is when a crisis hits, but this paradox suggests that beneath all the sophisticated reports and gadgetry, the traditional estimates of the risks you run when crisis hits are poor. There is mounting evidence that this is the case.

The paradox arises because of three aspects of investor behaviour. Following current finance theory, investor's view past returns, volatility, and correlations as independent from whatever they do. Investors use similar information, similar technologies and similar preferences and so end up hunting a similar group of markets. And concentrated purchases or sales of markets change their statistical properties - their likely return, volatility, and covariance. Understanding this process explains why the proposed solutions to mistakes in risk management cannot work.

What would have a better chance of working is the assumption that market covariances are part structural and part cyclical, and that in the future the cyclical component will act in the opposite direction than it did in the past.

Modern finance is an elegant description of a world that does not exist - it is an eighteenth-century world of ordered independence when instead we live in a world of strategic behaviour. Joining a crowded hunt for that portfolio that had the right balance of risk and return in the past, in the hope that it will deliver the same in the future, is not futile; it is dangerous. In finance, the future does not look like the past; it is a reaction to it.