Risk, theory, reflection: Limitations of the stochastic model of uncertainty in financial risk analysis

Barry du Toit

Principal Quantitative Analyst

RiskWorX

June 2004

---

1. Introduction

This paper argues that there is an important difference between the true uncertainty of equity returns, and the way in which modern financial theory models that uncertainty. In this paper we will examine the limitations of what I will call the stochastic (i.e. probabilistic) model of uncertainty[i]. The stochastic model is a crucial part of modern financial theory, and modern financial risk analysis in particular. By using a limited model of uncertainty, in which stock returns, although not predictable, can be described in terms of two key parameters (one stochastic), the stochastic model makes a complex problem mathematically tractable, and provides a means of domesticating the wildness of uncertainty. We will examine the limitations of this model, noting that its applicability varies across different topics in finance. We will be particularly concerned with areas where the assumptions of the model are especially problematic, but where the model is nevertheless freely used.

I will argue that modern financial risk analysis is institutionally biased towards forgetting the distinction between the stochastic model and true uncertainty. This leads to the uncritical use of statistical modelling in areas where it is not entirely appropriate, or, more generally, to the use of statistical modelling in an uncritical way. It is not the use of the models as such which is problematic, but rather their unthinking use. This leads to a systematic downgrading of the importance of investigative and reflective thinking in risk analysis. I will end this paper by restating the importance of such thinking. Reflection and theoretical analysis are the antidote to the unthinking use of models, and financial risk analysis needs to build these activities more explicitly into its institutions, practices and methodologies.

After briefly stating the problem of uncertainty in section 2 below, I move on, in section 3, to an examination of these issues not just as abstract debates about a "correct" model of risk, but as a developmental process typical of much of scientific development. The stochastic model of uncertainty has enabled the development of the flourishing discipline of financial risk analysis. This discipline arises from a set of initiating simplifications - the founding paradigm – as is the case with all scientific development. I will argue that the explanatory power of the stochastic model (sometimes genuine, sometimes illusory) has substantially influenced the concepts and practice of modern financial theory. As a result of this power it has come to dominate and shape the style of thinking in the discipline. This has created theoretical and practical shortcomings which need urgent attention. I highlight these shortcomings in section 4, and suggest a more integrated approach in section 5.

2. Investment risk and uncertainty

As we will see later, risk analysis is essentially casuistical: what constitutes a good answer depends as much on the specific features of the case at hand as it does on general considerations. So let me set up an example of the kind of financial risk with which I will be concerned with in this paper. Consider a typical pension fund portfolio consisting of a mix of domestic and local assets, such as shares, bonds and property. We want to know how risky it is to hold that portfolio. To answer that question we need to define risk. Let us say the questioner is a 30-year old women who wants to use the portfolio to buy a pension at the age of 65. She hopes the value of the portfolio will appreciate to buy her a better quality of life in retirement than her current savings level would lead her to expect. Her risk is that it will deliver a significantly lower quality of life, perhaps even unpleasantly so. How big is that risk, and should she take it? It's important to set the question up in this way, because the value of our concepts and models depends on the extent to which they allow us to provide better answers to these sorts of questions.

My background assumption here is that the future is ferociously difficult to forecast, and that the honest answer to the question of what the future holds is that we don't know. I am not going to argue this position at length here, but let's set out a few key features. The world is a complex place with very powerful forces evolving and interacting in complex ways. In this century we have seen a period of relative stability from the end of the second world war to the present. In that period, in most industrialised countries (particularly in the Western industrialised economies, but also in the Soviet block and in Asia), increases in national and personal wealth and quality of life were remarkable, and probably unprecedented in human history. But we need to be careful in extrapolating from this sample. From a historical point of view, we must acknowledge that in many ways the twentieth century constitutes a single-case sample from the history of human fortunes, and an extreme case at that. The conditions which made it possible may not persist. Our quality of life going forward may be radically altered by all sorts of developments, and the same applies to the value of the assets we invest in to provide for the future. That is to state things in perhaps the gloomiest of ways, but we need to be aware of that.

The point here, however, is not to retreat from the challenge, but only to be careful of underestimating it. We need to make whatever intelligent guesses we can about the future, and build those into our investment plans. And in the area of risk analysis and management we need to incorporate both those predictions for the future for which we have reasonable grounds, including those aspects of future uncertainty which we can model, as well as cater for the uncertainties which we cannot model at all. In fact we do this sort of thing in many areas of life all the time (for example in the sphere of medicine and health). We certainly cannot avoid the problem. Everyone has a de facto asset allocation scheme, even if everything they have is in a bank current account (or indeed in an overdraft). So we don't have to arrive at the correct asset allocation scheme, just a better one than most people currently have, and one based on plausible judgements. We have the tools already available for this. At the same time, we also need to be sure that we are not blinded by the dominance of some of our more succesful risk technologies.

My argument is simply that we do not plan for the future as well as we can because our understanding of financial risk has been distorted by the phenomenal success of a particular model of risk. The areas where that model is most appropriate have prospered, and the areas where it is least useful have either been neglected, or else have simply been approached using the conventional methods, regardless of relevance. We now turn to look at the development of this model.

3. The development of modern risk analysis: some Kuhnian insights

3.1 How paradigms develop

In this section I provide a description of the development of modern risk analyis which emphasises the social and practical dimensions thereof. This description of the development of a science is derived from the work of Thomas Kuhn, in which he emphasises the role of paradigms and paradigm shifts in scientific progress. Kuhn's work is often associated with a radical, relativistic concept of the history of science, but I won't be using the radical versions of those ideas here. Instead all I need is a particular story of how a science might develop. I don't even need to claim that this is a universal story: just that it is an interesting one which gives us some important insights, which I believe are useful in making sense of what is going on today in financial risk analysis. I identify 4 stages in this weakened version of Kuhn’s story.

i. Inauguration. In the beginning there occurs some sort of paradigmatic event. This might take the form of a practical event, such as a particular experiment, or it might consist of the introduction of a new concept or even a way of measuring. In the history of modern risk analysis I am going to point to two paradigmatic events: firstly, Markowitz's notion of portfolio diversification (here the paradigmatic event is Markowitz's seminal paper "Portfolio selection", published in 1952, but only really taken up in portfolio theory in the 1970's and after), and, secondly, the Black-Scholes derivative pricing model, in particular as the theoretical basis for the establishment and rapid growth of the modern derivatives markets. Embedded in these two events, and piggy-backing on their influence and power, are particular definitions of risk which, as we shall see, combine to give us the modern conceptualisation of risk.

ii. Vigour. We only identify events such as Markowitz's paper and the Black-Scholes model as paradigmatic in retrospect, because of the vigour of the new disciplines to which they give rise. The crucial point here is that it is not some purely epistemological criterion (truth, accuracy) which causes paradigms to succeed, but rather their ability to foster powerful and succesful research programmes. Successful paradigms allow us to analyse the world in ways which are both practically and theoretically fruitful, and both the above events have done that to an extraordinary degree. They introduce ways of making risk mathematically tractable, and allow the creation of a range of theoretically powerful analytic tools. At the same time, both find immediate practical applications in pressing issues of the day, providing solutions that are manifestly superior to earlier formulations.

iii. Dominance. Because of their practical success such paradigms come to dominate scientific and indeed intellectual activity in their area. Part of this dominance is appropriate, and represents a continuation of the vigour of stage two. But it isn't all good. One problem is that very dominant paradigms may act to stifle the development of alternative approaches. Another is that the dominant paradigm eventually moves beyond its areas of applicability and starts to extend into areas where the limitations of the inaugurating simplifications begin to be exposed. But it can take a while before those limits are recognized.

***iv. Revolution/Evolution*.** Revolution is of course Kuhn's core notion, but again I only need a watered-down version here. Kuhn tells a radical story of how a dominant paradigm finally collapses under internal and external pressures, to be replaced by an entirely new perspective. Often the trigger event is that the new paradigm proves successful in solving precisely those problems against which the old paradigm finally ran aground. In Kuhn's radical version, the old and new paradigms may be incommensurable, meaning that there is no common theoretical or observational language which can be used to compare one approach with the other, with the result that the adoption of the new paradigm cannot be justified on rational grounds. Some sort of "leap of faith" will then be necessary. As the history of science shows, though, it is only in some very extreme cases that something so fundamental occurs. More generally new approaches may incorporate the old as a special case, or just start to function alongside the old, perhaps focusing on different content. In this paper I will be suggesting some ways in which modern risk analysis might profitably evolve.

3.2. Two paradigmatic events

3.2.1 Markowitz: Standard deviation as the measure of risk

Markowitz's work on portfolio diversification is a key point in the history of the development of this discipline, perhaps the foundational point. As with any science, progress requires the construction of concepts which make the subject matter tractable, and typically this means mathematical tractability. Such concepts take the messy confusion of the world and identify key features which can be measured and modeled. Successful models allow us to extend our knowledge and understanding into new areas, until eventually the weight of the initiating assumptions acts as a greater and greater drag on expansion, and a new model takes over, better able to deal with the problems on which the old model foundered.

Although Markowitz published his ideas in 1952, in what was pretty much their mature form, they only came to influence financial decision-making in the 1970's. Rather remarkably, the modern edifice of mathematical modelling, risk analysis and financial engineering only really began in that decade, although many of the important ideas had been around for years, and had in fact been taken to relatively advanced levels in some other statistical and economic sciences. Many reasons have been put forward for the timing of the rise of modern financial engineering, including the volatility in global markets in the 1970's, the rise of inflation in stagnating economies, the growth in computational power and the spread of access to computational technology, and so on. For whatever reason the time was right, and Markowitz's ideas provided the necessary conceptual framework.

Markowitz's key idea was, of course, the notion of portfolio diversification. His work pointed out the existence of one of the very few genuinely free lunches in finance. Here is an example. Suppose you have two shares, both of which have the same expected return (say 5% per annum) over the investment time horizon (say 4 years). Suppose both shares are risky - there is no guarantee that the return in any given year will be 5%. In some years it may be more, in some years less, and the degree of uncertainty is the same for each share. But suppose, in addition, that the two shares are negatively correlated - in years that share A does well, share B does badly, and vice versa. Then the portfolio consisting of a combination of the two shares will have the same expected return as the shares themselves, but less risk. The reduction in risk is free in the sense that no compensating reduction in expected return is required.

Figure 1: Combining two shares produces a portfolio with lower risk and the same expected return

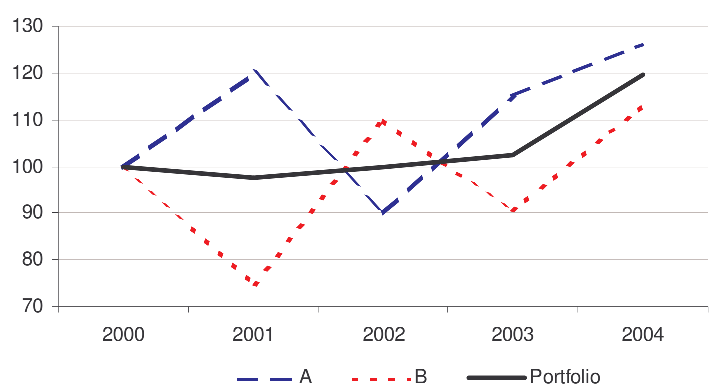

Markowitz's notion of diversification gave rise to much of the current practice of portfolio optimisation (the construction of portfolios which maximise the ratio of expected return to risk), and much of the discipline of financial engineering. Indeed in some ways he can be said to have invented the idea of financial engineering , showing how to construct objects (portfolios) out of raw materials (shares, bonds, cash, etc.) in such a way as to maximize the desirable properties of the construction, and largely taking the atomic properties of the objects (returns, volatilities, correlations) as given. Markowitz also emphasised the issue of risk aversion: that human beings, given two investments of equivalent expected return, will choose the less volatile alternative, or, more generally, that there must be some increase in the expected return in order to compensate the risk-averse investor for some increase in risk. These concepts were central to the paradigmatic nature of Markowitz's work, setting out a clear programme for applied risk analysis in portfolio construction. But the concept which is most important for us here is one which is in fact of only peripheral significance in Markowitz's system: the choice of standard deviation as the measure of risk. Markowitz did not spend a lot of effort on his choice of risk measure, and indeed his results mostly hold if other measures of risk are substituted for standard deviation. But it was the definition he chose and, riding on the back of the enormous influence of his ideas, it rapidly became the standard definition of risk in financial risk analysis.

Markowitz essentially gave a formal expression to the notion that investment risk is best characterised as the variability of the expected returns, and then linked that general idea to the specific statistical concept of standard deviation (or variance, which is just the square of the standard deviation). At this point we can note that variability is already a simplification of what might count as investment risk. But the concept is so broad and accommodating to a range of risk measures that I am going to ignore it at this point. Rather we will focus on the specific identification of risk with standard deviation, which was, until the introduction of VaR methodologies[ii], the primary definition of risk in investment theory and practice.

Let's note carefully exactly what the definition of risk as the variability of expected returns means. First, this concept is entirely unobservable, and refers to our present views about the future performance of the share. We think the share will return 20% over the next year, but we know that is not certain. All sorts of global and local factors, as well as the unpredictable response of the market to changes in these factors, will determine the actual return. The range of possible returns may be quite large. It is the range and distribution of these alternative outcomes which determine how risky we believe the share to be. We could theoretically attempt to model this uncertainty directly, by devising a model of the sensitivity of the share returns to various factors, and pushing this modelling down to some level where we can quantify in some meaningful way the range and probabilities of the future states of the variables. But this is currently just a fantasy, and may well stay that way permanently. So far, then, the definition does not do all that much for the key attribute of tractability. Two more steps are required.

The first step is to define variability (or at least the aspect of variability that matters) as the standard deviation of returns, which Markowitz does. Standard deviation is defined as σ (sigma), where:

$$\sigma = \sqrt{[(\sum (x_i-\mu)^2)/n]}$$

with

x_i = the i^{th} possible outcome

μ = the arithmetic mean of the outcomes, and

n = the total number of possible outcomes.

Note that this is something of a conceptual leap of faith. It isn't obvious, firstly, that we can properly model our understanding of risk with a single measure of risk. And if we can, then it isn't obvious that standard deviation is that measure (there are many other possibilities). For example, do we want our measure of risk to be based on volatility relative to the mean return (x_i-μ), or should it simply use the non-adjusted, absolute x_i? And do we want our measure of risk to utilise the square of the individual deviations $(x_i-\mu)^2$, rather than just, say, the average of the absolute value of the deviations abs(x_i-μ)? And should our measure of risk weigh upside risk as much as downside, as the standard deviation does?

The point is that Markowitz's embedding of standard deviation in his model of portfolio construction, which was to dominate the new discipline of financial engineering in the 70's, performs the classic function of an inaugurating simplification in the establishment of a paradigm. Those who followed on from Markowitz were able to move forward without first having to deal with the difficult and time-consuming issue of how best to characterise risk, but could simply adopt a ready-made and explicitly mathematical definition. Because of the tractability now bestowed on the matter of risk, and because the definition was reasonably well-suited to its initial areas of application, the science of financial engineering was able to move forward with impressive vigour. In direct line of descent from Markowitz's idea, a complex and sophisticated science of portfolio optimisation was created, able to draw on a range of resources in matrix algebra and other mathematical sub-disciplines, and Markowitz's original idea, supplemented with additional theoretical assumptions of varying plausibility, gave rise to a range of important ideas such as the Capital Asset Pricing Model, the associated concept of beta, notions of risk-adjusted returns, and many others, as well as providing a precise language for the grand theory of the efficient market, which dominated academic debate over the last quarter of the twentieth century.

So far we have discussed what might be called the theoretical aspect of the tractability of the standard deviation measure. One final step, required for practical tractability, was to make the concept measurable in practice. Recall that the concept is essentially intended as a measure of the variance of the *expected* returns. There is no way of observing this. The final step, then, is the simple assumption that the variance of the expected future returns will reflect the variance of returns in the past. This is really quite a big step, but it does make the standard deviation easy to measure, and there isn't really any plausible alternative. Suppose we take as our universe the last 1000 days of returns on some share A. Then we simply assume that the probability of any return occuring tomorrow is just the frequency with which it occurred in the past. The standard deviation formula is then simply applied to the historical data: the expected return becomes the mean return over the historical sample, and n is the number of days in the sample.

In this section I have shown how Markowitz turned what was in fact a complex and intractable issue into something which could form the basis of a progressive science. What made his choice of risk measure important was not the force of his arguments in favour (in fact he offers hardly any), but rather the utility of the measure (as well as simply the specification of *some* measure), as one of the basic building blocks (along with return) of the science of financial engineering. The crucial elements contributing to this utility were the mathematical form of the definition, the mathematical tractability of the measure itself (as opposed to, say, the mean absolute volatility, which is algebraically awkward), and the fact that it (or at least an adequate proxy) can be empirically estimated.

But the measure is still a little anaemic when it comes to the serious work of building an intellectual empire. It has nothing to say about the nature of the return-generating process itself, a limitation which seriously restricts the extendability of the model, especially across time horizons. Let's go back to our example: how risky is an investment in a given set of shares over a 10-year time horizon. If we attempt to apply the standard procedure described above we would need a sizeable set of 10-year horizon data on the portfolio in question. Typically we don't have anything like this much data, and the fix used is to calculate the standard deviation for a 1-year horizon (or even shorter), and then extrapolate this result to the 10-year horizon. But this requires additional assumptions about the way in which stock volatility behaves. Markowitz provided an essentially passive quantitative description of variability. To move between different horizons we need some assumptions about the dynamics of the process which gives rise to variability. We look at the historical solution to this in the following section.

3.2.2 Black-Scholes: Geometric Brownian motion as the model of risk

I will show here how the introduction of the Black-Scholes formula for pricing options, and its phenomenal success in creating the modern world of derivatives (what Bernstein calls the "fantastic system of side-bets"), helped to legitimate the next key step in the taming of risk, namely the transformation of volatility as a descriptive measure of distributions, agnostic on the generating process concerned, to a parameter at the heart of the generating process itself. Let's see how this works.

The Black-Scholes formula provided the key breakthrough in showing how to value option positions, such as calls or puts. Let us consider a simple call option on a share. The share is currently trading at R90, and the option gives us the right (but not the obligation) to buy the share in 1 year's time for R100. If the share is trading at R120 in 1 year's time, we exercise the option and take our R20 payoff. If the share is trading at R80, we do not exercise and simply lose whatever we payed for the option in the first place. More formally, then, we have a 1-year call option with strike K. Let the price of the share be S. Then in 1 year's time the payoff from the call option will be max (S_1 - K,0), where S_1 is the value of the share at the end of the year. The "max" is the key to the essence of optionality. The option only pays out when it is in the money. If the share price ends below the strike price, the option expires worthless, and you lose whatever you payed for it, but no more.

Figure 2: Payoff profile of a call option with strike 100.

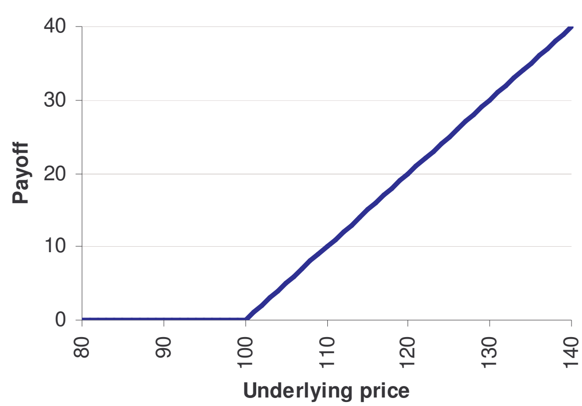

The question is how much should you pay for such an option? Rather extraordinarily, Black and Scholes demonstrated that what you should be prepared to pay for an option is independent of your view of how well the underlying share or index is going to do. It doesn't matter at all whether you expect the share to rise or plummet over the next year: the price at which the option will trade in the market (given certain necessary conditions) will be insensitive to that (seemingly critical) issue. They prove this by showing how to construct a portfolio which is long 1 option and short a certain quantity (called the *delta*) of the underlying share itself, which portfolio is entirely *riskless* for the next instant in time. (In the case of a call option you would be short the underlying share, in the case of a put option you would be long). This portfolio is riskless because it removes the question of the expected return of the share from the equation. We don't know the expected return, but it turns out it doesn't matter because in the combined portfolio the effect of the expected return on the option's value is exactly cancelled out by its effect on the share position's value. Because the portfolio is riskless, it must grow at the risk-free rate (normally assumed to be the rate associated with Government bonds), otherwise there would be arbitrage opportunities in the market (that is, you could borrow money and invest it in the option/delta combination and make a riskless profit). And equally, because it is riskless, the appropriate discount rate (the rate at which you work out what the expected profit at maturity should be worth today) is also the risk free rate.

Note that it is the action of the arbitrageurs which enforces the Black-Scholes price on the market. Suppose you are so sure that the price of the underlying will go up that you are prepared to buy very large quantities of the call option. With a non-derivative instrument, such buying pressure would push the price upwards. But for a derivative instrument, the moment you bid above the Black-Scholes price an arbitrageur will be there to sell to you, thus cancelling the pricing pressure of your demand. And the reason the arbitrageur will sell (in theory with no limit) is that he can delta-hedge the option according to theBlack-Scholes formula and make a riskless profit.

The insight is absolutely fundamental, and worthy of the Noble Prize it eventually received. Apart from providing the foundation for the whole of the derivatives industry, it also gave rise to the important theoretical conceptions of no-arbitrage and risk-neutral pricing.

Even more, Black and Scholes were, on the basis of this insight (and a little mathematics), able to derive an exact mathematical formula for the price of any option. The details of the pricing formula need not concern ourselves here, except to note that the formula requires four inputs: the term to maturity, the strike level of the option, the current level of the underlying instrument, the risk free rate, and the volatility (the standard deviation) of the underlying instrument. Note that there is no reference to the expected return of the underlying. The first three are generally easy to obtain. The key input is then an estimation of the expected volatility. Since Black and Scholes have shown we can ignore the expected return, volatility remains as the only significant unknown factor in option pricing. The reason volatility matters is essentially because the payoff profile is asymmetric (that is, negative values of S-X all pay off as 0, whereas positive values pay out their actual amount). As volatility increases, the increasingly high positive values of S-X add value to the option which is not cancelled out by the equivalent negative values, which remain with pay-off 0.

Knowing the price of the option and knowing how to hedge the option made it possible for banks to confidently write (i.e. sell) options to anyone who needed them, such as fund managers wishing to protect their investments. The confidence inspired by the Black-Scholes equation thus created the conditions for the massive growth of the modern derivative markets.

Riding on the coat-tails of this success was the exact model of the underlying used by Black and Scholes. In order to force their mathematical equations through, Black and Scholes required a precise mathematical model for the way in which the underlying instrument evolved through time. The formula they used has also become a central pillar of modern risk analysis, and it embodies what I call the stochastic model of uncertainty. Recall our earlier comments about the deep unpredictability of equity returns. Black and Scholes use a model which appears to tame this wildness. Note that this a model of the underlying instrument, not of the hedged portfolio, so there is a genuine wildness to be tamed. The model of the underlying process used by Black and Scholes, now ubiquitous in finance, looks like this:

$$r_t = \mu * \Delta t + \sigma * \varepsilon * \sqrt{\Delta t}$$

This says that the return r, at any step t, of a financial instrument is made up of two components. The first is the expected return, μ, which is fixed, multiplied by the size of the time-step Δt. The second component is a random drawing ε from a standardised normal distribution, scaled by the volatility of the instrument, and by the square root of the time-step. So let's say that we are looking for the 1-year return on an instrument which has an expected return of 15% and a volatility of 20%. With a standardised normal distribution, the mean, median and mode are all zero, so the most common random drawing would be 0, giving:

r = 15% + (20% * 0) = 15%, the expected return.

But in another world, or another year, the random drawing might be -1, giving:

r = 15% + (20% * -1) = -5%, and so on.

Because we are drawing from a normal distribution, the shape of the distribution resulting from thousands of repeated drawings will be the well-known bell shape.

The exact form of this process posited by Black and Scholes is known as geometric Brownian movement, which, in addition to the features I have mentioned, has the property that the random drawings are independent. Simply put, the probability of a return of, say, -10%, is always exactly the same, even if, say, the previous day's return was also -10%. This contradicts most people's intuitions about the market, and there are adjustments to the model which attempt to incorporate such autodependencies.

This formula represents the high point in the taming of uncertainty in modern finance. It takes the notion of risk as the variance in returns introduced by Markowitz, but goes a step further. In the Markowitz world the standard deviation of returns is simply a passive statistical description of a set of states of the world: either of the possible states in future-expected space, or the set of actual recorded states in historical space. In other words, in the Markowitz system we don't assume we know how the future or past distributions are generated, just that we have a useful statistical tool for describing the results of the process. After Black and Scholes, the standard deviation notion of risk is embedded in the heart of the return-generating process itself. The gains in terms of mathematical modelling are immense. But the messy uncertainty of forecasting future returns, and the earlier questions about the best way of understanding risk, are easily forgotten.

4. Limitations of the stochastic model

Modern risk analysis has cut an impressive swathe through the complexities of risk, culminating in the stochastic process model. Such processes generate series of returns, or prices, which look like the price series of typical shares across their real histories, but with the useful difference that we have a precise model of how the returns are in fact generated. But we need to remain aware of what has been sacrificed in order to achieve that precision.

Essentially, the stochastic model substitutes a "tamed" version of uncertainty for the "wild" uncertainty of the real world. I raised this distinction in general terms earlier in this paper, and now we can make this distinction more precisely. While the outcomes under the stochastic model cannot be predicted, and are thus uncertain, the risk process has been parameterised, and the values of the parameters (the mean return, and the stochastic parameter of standard deviation, or volatility) are known. The notion of uncertainty has changed. A generalised uncertainty has been replaced by a model in which uncertainty is contained and parameterized: the real uncertainty of share returns is split neatly into two parts, one of which is deterministic (although only ever a forecast) – the expected return term – and the other of which contains the uncertainty, interpreted as the volatility of a stochastic process. Note that the uncertainty in the estimation of the volatility parameter itself is largely ignored, both in theory and in practice.

This tamed version of uncertainty represents an important resolution to the problem of uncertainty. The history of human responses to uncertainty has tended to run to two extremes: on the one hand some form of fatalism in the face of an entirely unpredictable future, and on the other the attempt to remove uncertainty by discovering the scientific laws governing the process concerned. The stochastic model provides a third way, which acknowledges uncertainty and unpredictability, but still allows for the creation of a powerful applied discipline - financial engineering - in which risk can be modeled, compartmentalised, mitigated, hedged and insured against in order to meet our requirements and assuage our fears. This gives us mastery of a sort over even those parts of the world where deterministic solutions cannot succeed. It is an important mastery, but it is also limited and at times misleading.

If we mistake map for reality, of course, then we will not notice any inadequacy on the part of our models at all, so let's be clear about the extent to which the stochastic model is indeed a model. If the daily returns of shares were generated by some process which was structurally similar to the geometric Brownian motion process posited by Black and Scholes, then there would not be any problem. The only difference between model and reality would then be that of noise, and we would be faced with the orthodox statistical problem of estimating population parameters from sample statistics. If this was true then there would always be an objective, albeit unobservable, process volatility, which we could try to estimate from the incomplete sampling we have available.

But the uncertainty of financial markets almost certainly does not arise from such a mechanism. Process volatility is not just unobservable, it does not exist, except perhaps in weakened form as a summary measure or a market propensity. In addition to replacing a generalised notion of uncertainty with a more structured one, the stochastic model uses a model of the underlying process which is an inadequate likeness of the real thing. The behaviour of financial prices is the result of the extremely complex interaction of a variety of processes – essentially, the processes which supply new information to the market, and the way in which the heterogenous mass of market participants react to that information. Markets are, of course, precisely mechanisms for bringing a multitude of players and processes together into one single bit of information, the price of share A at time t. This price-formation process is almost certainly non-linear and historical (that is, the relationships evolve through time). Sometimes the mechanism can indeed be transparent and simple: new information such as that of the 9/11 attacks on the World Trade Centre dominated all other news on that and subsequent days, and was interpreted in very similar ways by most market participants. So the broad mechanism governing market pricing over that period was fairly straightforward. But most of the time, of course, prices arise from the interaction of a multitude of relatively evenly-balanced factors.

In fact it is probably true to say that the process which generates market prices is fundamentally deterministic, rather than fundamentally stochastic. That is, the uncertainty is not a fundamental attribute of the generating process, as it may be, for example, in quantum mechanics. The uncertainty of financial market pricing arises from the fact that the massive complexity of the process dwarfs our intellectual capacity, and we therefore have neither understanding of, nor the power of prediction over, market prices. (Of course, the possibility of predictive power is ruled out not just by complexity but also by the reflexive nature of our knowledge of markets, in that such knowledge in turn changes the nature of the very thing we seek to know.) But although prediction with any sort of certainty will always remain an impossible goal, that does not mean that the generating processes are stochastic. And, even more strongly, they are not the sort of simplistic stochastic process utilised in financial modelling as we have described it. Put another way, the stochastic approach makes the assumption that the underlying process is governed by statistical laws, rather than the weaker assumption that the outputs of the process can be adequately summarised by some appropriate statistical measures.

Thus far I have presented a general argument to the effect that the stochastic model is only a model, and not even especially homologous to the reality it seeks to describe. But, of course, it shares, by construction, the property that it can be described in terms of the two parameters (or moments) by which we characteristic the distribution of market returns, namely the expected return and the variance, and that is of course primarily what we are concerned with. From a practical point of view, then, it may still be an effective model, and indeed, as I have stressed, this paper in no way seeks to deny that. But it is also worth pointing out that even from a pragmatic point of view there are some important weaknesses in the stochastic model.

We can examine these shortcomings by looking initially at where the model succeeds - where there is a good fit between the model and the content. One case is that of derivative hedging, where derivatives traders, looking to hedge a derivative position (i.e. construct a riskless portfolio of the derivative and the delta quantity of the underlying share), make use of the stochastic model with a high degree of success and with few grounds for concerns (at least in terms of the issues we are concerned with here). Another area would be in the estimation of short-term (1-day) risk in markets such as the global currency markets. Let's note two common properties of these areas which make the fit such a good one.

Firstly, the expected return component of the stochastic equation is either irrelevant or negligible. In the case of the derivative hedger, the expected return component of the equation can be ignored, precisely because the overall portfolio is neutral with regard to the expected return. In the case of short-term currency risk movement, the expected return is not irrelevant but it is negligible, because the day-to-day volatility dwarfs the drift term in the equation. This is extremely useful because forecasting the expected return is a lot more difficult than forecasting future volatility. So the model is most at home in situations where the mean return is insignificant. In forecasting long-term equity risk (actually anything from a few months forward to a few decades) the mean return term becomes increasingly important, and the value of the model starts to decline.

Secondly, the investment horizon is short. Derivative hedgers typically rebalance their delta on a daily basis, or even more frequently in very volatile markets. In the case of currency trading the risk management process works on very short-term horizons, so the risk analysis does not need to extend far beyond a forecast for the next day. One advantage of a short horizon is simply that the expected return component is likely to be less important, as discussed above. But a second, independent, advantage is that the estimation of the necessary parameter - standard deviation - is likely to be much more accurate, simply because we have access to a much larger sample of holding period returns. If we are concerned with longer investment horizons - say a month or a year – we typically do not have sufficient good quality data on which to base our forecasts. Analysts often deal with this by using shorter return periods as their basic data – daily returns, for example – and then scaling the results up to the required horizon. But any choice of scaling method assumes a relationship between the short-term and long-term dynamics of the share process, and of course it is precisely the long-term dynamics that we have not been able to sample adequately.

In this section I have attempted to show that the stochastic model of uncertainty is powerful but limited. In some areas the limitations are less relevant, and we can use the model with confidence. But in other areas the limitations are substantial, and we need to be aware of them. In the final section I discuss some ways of dealing with this issue.

5. Theory and reflection

Risk analysis will continue to evolve, along two complementary paths. The first path will be the continuing enhancement of the risk-analytic toolkit – the set of modelling and measurement options which has developed so rapidly over the last few decades. But possessing the most powerful tools is not sufficient. The generation of useful risk measures and risk management procedures requires, in addition to the technically most advanced toolkit, a proper understanding of the nature (and limitations) of the tools, and, most importantly, an understanding of the nature of the goal. In other words, what is also required is the intelligent and thoughtful application of the tools. In this final section I argue that in addition to the development of its technological reperoire, risk analysis also needs to develop its theoretical and conceptual resources. This means, for example, analysing what counts as a successful risk measure. Decisions about the most appropriate form of modelling and the most appropriate combination of risk statistics flow from this understanding.

Before saying more about the significance of theoretical analysis, I will briefly look at some developments in what I have called the technological resources.

5.1. Enhanced technology

The stochastic model is the best mathematical model of financial asset dynamics we have, and it will continue to develop as a way of modelling risk. Many of those developments will make the model fit increasingly well with the real world of risk, and extend its applications into areas where it now seems less capable.

One particularly fruitful line, for example, has been the development of increasingly complex models of the nature of volatility. The original Black-Scholes equations assume a constant process volatility over the life of the option. But practitioners and theorists quickly realised that process volatility is likely to be an evolving parameter, rather than a static one, and we have accordingly seen the development of increasingly complex models of the dynamic character of the volatility parameter, such as the use of stochastic volatility models, where not just the share return process but the volatility parameter itself are taken to be evolving stochastic processes. These theories extend our practical and theoretical resources, and continue to deepen our appreciation of the nature of financial risk.

In addition, there are alternative approaches to modelling which are less dependent on a set of underlying assumptions than the full-blown stochastic model. One simple but important alternative to strong parametric modelling, such as in the classic stochastic model, is simply to look at the unadjusted historical data, not filtered through any sort of model of volatility. For example, if our concern is with the riskiness of a ten-year allocation to equities, we might look at the history of actual ten-year returns to the relevant index. The probability distribution of these returns will tell us some interesting things about the type of risk we face going forward. This kind of approach – raw historical simulation – has a number of serious weaknesses. For example we seldom have enough data if we are concerned with longer-term returns. Secondly, the global economic, social and political context is in constant flux, has changed dramatically over time, and will continue to do so. Just because something has not happened in the past does not mean it may not happen in the future, and, conversely, issues which arose in the past may no longer be of any significance now or for the foreseeable future. Nonetheless, weak modelling methods such as historical simulation are a useful corrective to the naïve use of strong modelling.

So one way of enhancing our understanding of risk is to ensure that we use a variety of modelling approaches – parametric and non-parametric, historical and analytic, strong and weak. This can be very effective when the models complement each other. For example, a parametric model allows one to manipulate the inputs and test a wide variety of hypotheses, whereas a historical simulation provides an important reality check. But the use of a number of different approaches points us immediately to a range of questions about how we go about combining them: which do we use, and when, and what do we do if different approaches point in very different directions, and so on?

It's clear that combining different approaches requires a clarity of purpose and some sort of appropriate theoretical framework. This leads us towards the final point I want to make here, which concerns the importance of enhancing our theoretical analysis of risk.

5.2 Enhanced analysis

Analysis is the antidote to the unthinking use of models, and it is the unthinking use of models, and not the models themselves, which is the target of our criticism here.

Applying one's mind

Part of what is required is simply described as applying one's mind, or developing a more reflective approach to risk analysis. The dangers of the mindless application of risk technology have been raised, with increasing urgency, by a number of writers, of whom Taleb is perhaps the most strident. What is important to note is that this unreflective analysis does not result just from cognitive inadequacies on the part of the risk analyst. Rather it results from the combination of a number of precipitating factors in the typical institutional environment. Thus in many cases the risk manager is overburdened with the daily effort of generating plausible risk numbers, with little time or energy left over to reflect on their meaning. Similarly, quantitative analysts prefer to deal with questions which they know how to answer, and which fall into their area of technical expertise. They accordingly prefer to be presented with a quantified measure of risk, the value of which they are able to calculate and estimate, rather than to reflect on the suitability of the measure itself. And there are similar influences on the wider set of interested parties as well. Consultants and advisers, for example, like to show that they are applying some sort of systematic or scientific method, not just voicing (subjective) opinions. And investors themselves are frequently reassured by the appearance of quantitative rigour. Lastly, we should note that, as we have seen in this paper, the language of financial risk analysis is dominated by some very powerful quantitative paradigms. And all of this, of course, gets wrapped up in the deadening inertia of "best practice".

Taleb's altimeter analogy is useful here. In his critique of VaR methodologies in risk analysis Taleb compares the employment of VaR with the use of an altimeter in an aeroplane, and more especially, an altimeter which is right only 95% of the time (VaR models forecast the worst case outcome at a given level of confidence, typically 95%). Of course, when the altimeter does get it wrong the plane will crash. The analogy is not exact, since banks use VaR numbers in such a way that loss events beyond VaR will not result in disaster, but that is irrelevant to Taleb's point, His point is that the danger inherent in altimeter usage is that its *unquestioning* use will stop pilots from using a range of more intuitive risk measures, such as looking out of the window! The real danger of using high-tech quantitative risk measures lies not in their inaccuracy or lack of validity (although Taleb has serious reservations here as well), but in the degree to which they are used as a substitute for thinking, and generate a false and dangerous sense of security.

So it is important that someone, somewhere, applies his or her mind to the problem of risk. The main point here is institutional: we need to create the time and space within our instutions for the reflective exploration of risk, we need to assign responsibility for this function, and we need to reward it. In addition we can restructure the delivery of risk data in such a way as to facilitate exploratory thinking on risk. One simple technique is to provide a richer set of data to the risk manager, together with an appropriate set of exploratory tools. This is at odds with the usual requirement of a specific set of measures whose production can be automated and to which a set of straightforward risk limits can be applied. However it is not a replacement but rather an important complement to such an approach.

Applied theoretical analysis

But the analytic step consists of more than just the injunction to apply one's mind. It also consists of prioritising and foregrounding the stage of theoretical analysis.

Its useful in this reqard to identify 3 stages in the risk management process:

What constitutes risk in this contex?

How best can we measure and monitor it?

How best can we manage it?

The first step is the step of theoretical or conceptual analysis. In theory this step could be trivial. Suppose for some investor standard deviation is in fact a complete definition of what counts as risk for them. An easy example would be that of a fund manager who is rewarded for his or her risk-adusted return where risk is defined as the standard deviation of returns. In such a case the conceptual stage collapses. By construction, some quantitative measure of risk just is precisely what is meant by risk in this context. So here we can proceed directly to the technical questions of how best to model the forecast standard deviation of the portfolio, and from there to the question of management. In this case risk analysis is almost entirely a technical process.

But this is a contrived situation. The point is that it is not at all like this for almost all real investors, for whom any quantifiable measure of risk is only a proxy, highlighting one facet of the true risk which is theoretically complex, ambiguous and multi-faceted. We need to look very carefully at exactly what constitutes risk for a particular investor in a particular context. The answer needs to be both sufficiently rich to correctly capture all the key elements of the contextual risk, and clear and analytical enough to guide the construction of suitable risk measures and effective risk manageement procedures. For most investors, then, answering the conceptual question is both important and complex. Let's look in more detail at some reasons why.

Risk is conceptually complex. There just isn't any single way of capturing the notion of risk in any reasonably compact definition. Consider as an initial framework a formal definition of risk as the adverse subset of the set of outcomes of a particular action or process. Obviously we need to add quite a bit of substance to this definition to make it useful, but as soon as we attempt to do this the complexities begin to multiply. For example, just what constitutes an outcome. In the analysis of investment risk one component is the specification of the horizon of investment. And this frequently has no clear answer. The answer may be fuzzy: just how long is "the long term". Or it may not be known: "the rest of my life". Moreover, the notion of outcome may not even be a useful one: we may be just as concerned about risk along the way (interim risk) as the risk at some notional terminal point (Kritzman). In addition to the problem of horizon, we are also concerned with exactly how we describe an outcome. Is an outcome the amount of money lost or gained, or the probability of a loss exceeding some critical amount (ruin), or something else? And then there is the matter of what makes an outcome count as "adverse", and how we develop some appropriate metric of adversity. This issue is often characterised as ascribing a utility to all the possible outcomes, but this is not easy and, in addition, the notion of a utility "calculus" may be a formally inadequate way of describing the human response to risk.

Of course we don't have to come up with God's definition of risk, but only one suitable to provide some reference for risk management in the human world. But we should not forget that that process is complex, and that any operational definition will be nothing but a thin quantitative proxy for the richer underlying concept. Providing we remember that, we stand a chance of being able to use in the operational definition in an intelligent and useful way.

Risk is multi-faceted. It's also likely that, even once we have tied down the conceptual complexities, the definition of risk in the typical financial context is likely to be multi-faceted.

For example, you may be concerned both with investment risk at some terminal point, and with risk along the way (interim risk). You may be concerned with how likely a loss might be, but also with how big it will be when it does occur. You might be concerned with "typical" day-to-day volatility, but also with highly unlikely extreme events. Because risk typically has more than one facet, we may require a number of ways of measuring and managing risk, and, most importantly, some way of ordering our various concerns in guiding the management process. Once we enter this plurality of measures and models we require a clear theoretical framework to develop the appropriate decision rules. And this framework must emerge from a careful analysis of exactly what constitutes risk in this instance for a particular investor, or class of investors.

The theory of risk is deeply multi-disciplinary. The theory of risk is essentially multi-disciplinary, covering a range of disciplines including economics, psychology, statistics, mathematics and philosophy. As a scholarly discipline it does not have a clear home. In particular the mathematical nature of financial risk analysis may act as something of a barrier to integration with other modes of thinking about risk. Although difficult, it's clear that an integrated theory of risk is a highly desirable goal.

Applied risk analysis is strongly casuistical. Applied risk analysis is significantly casuistical. While the development of an overarching theory of risk is important and useful, risk analysis by its nature is always specific to the particular case at hand. We always have to work out in detail what risk means for any given investor in any given context.

Risk theory plays a guiding role in measurement and management Since Popper we have understood that the role of theory is crucial in the practice of the scientific method. In scientific research theory generates the testable predictions which are the means by which we can justify our beliefs. In risk analysis theory plays a crucial guiding role, without which our investigation is effectively blind. For example, only when we have a clear description of what matters as risk can we judge what risk measures will be appropriate. The choice between, say, VaR and standard deviation as a risk measure rests on number of criteria, including important technical criteria (such as the extent to which we can accurately and reliably model the measures), and the statistical properties of the measures themselves (such as coherence). But clearly the overriding concern is that the measure best fits our concerns about risk in this context. If we are concerned about the likelihood of large losses then VaR is the better measure, if our concern is with the typical, day-to-day volatility of the portfolio then the standard deviation may be more appropriate. But note that it is not as simple as saying that validity is all that matters and that other concerns are secondary: if, for example, VaR is clearly more valid and appropriate, but we can only measure it very poorly, then standard deviation may still be a better choice. Being able to make this decision depends on a deep understanding of exactly why VaR was more appropriate, and exactly what is being lost if we substitute another, more measurable, alternative. Developing a clear theoretical description of the contextualised risk gives us the only legitimate means of deciding on the best statistical proxy.

Finally, the theory also tells us how we go about applying the measures in the risk management process itself. Here understanding the utility issues in the definition of risk may be crucial. For example, there may be certain risk outcomes which are simply intolerable for the investor or for the portfolio manager, in that they would result , say, in insolvency or the failure to provide for basic medical costs. In such a case modelling may not even be necessary: since we cannot tolerate the possibility of such a loss, however small, there is not much to be gained by estimating that probability. We may, for example, require the employment of derivative structures to create a "floor" under the portfolio, ensuring that losses greater than a certain size simply cannot happen. And if such a measure is in place, that will change the way in which we model and manage the remaining risk in the portfolio. For example, with a derivative hedge in place, we may no longer be concerned about the ability of our model to capture extreme events, since these have already been taken care of. Rather, we may simply require the measure best at measuring typical, non-extreme volatility.

In fact it is worth noting that the theory of risk analysis is ultimately a pragmatics – that is, it is a theory whose purpose is not description as such, but guiding practice. Good risk measures are simply those which enable us to manage risk effectively – being accurate descriptions of risk or the potential risk is simply a means to that pragmatic end. In fact the traditional process model of risk analysis outlined earlier is misleading in this regard. We defined the stages as:

Define risk
Measure and monitor risk
Manage risk

Earlier we were concerned with highlighting the significance of the first stage. But now we can see as well that the first stage should be conducted with the goal – management – firmly in mind.

6. Conclusion

Market risk analysis requires a theoretical structure adequate to its complexity. Such a theory will always have to be more than just a technology. Risk is a complex topic at the intersection of statistical, economic, psychological and philosophical concepts. At the same time it is always an embedded topic, rooted in a particular investment context.

The complexity of risk arises from many sources. In this paper we have focused on highlighting the deep uncertainty of financial risk analysis, and emphasised that we should not be blinded to this by the success of the quantitative technologies that have emerged in the last few decades.

We have sketched one way in which modern risk analysis has remodelled uncertainty – the stochastic model of uncertainty. The uncritical concept of such reduced versions of risk may lead us to think that risk analysis is a purely technical exercise, and there are powerful institutional forces which promote this view. But risk analysis is not just a technical exercise. Risk analysis must always be rooted in a reflective and investigative thinking about risk. We need to continue to develop the institutional resources, practices and theoretical tools required for this task.

Bibliography

Bernstein, Peter L. 1996. *Against the gods: The remarkable story of risk.* John Wiley and Sons.

Gigerenzer, Gerd. 2002. *Reckoning with risk: Learning to live with uncertainty.* London: Penguin.

Hull, John C. 2000. *Options, futures and other derivatives.* 4^{th} Ed. Prentice-Hall.

Jorion, Phillipe. 1997. *Value at risk: The new benchmark for controlling market risk.* New York: McGraw-Hill.

Kahneman, Daniel, Slovic, Paul and Tversky, Amos (eds). 1982. *Judgement under uncertainty: Heuristics and biases.* Cambridge University Press.

Kritzman, Mark P. 2000. *Puzzles of finance: Six practical problems and their remarkable solutions.* New York: John Wiley and Sons, Inc.

Kuhn, Thomas. S. 1996. *The structure of scientific revolutions.* Chicago: Chicago University Press.

Osband, Kent. 2002. *Iceberg Risk: An adventure in portfolio theory.* New York and London: Texere.

Rahl, Leslie. 2003. *Hedge fund risk transparency: Unravelling the complex and controversial debate.* London: Risk Books.

Shiller, Robert J. 2003. *The new financial order: Risks in the 21st Century.* Princeton: Princeton University Press.

Siegel, Jeremy J. 1998. *Stocks for the long run.* 2nd ed. New York: McGraw Hill.

Taleb, Nassim Nicholas. 2001. *Fooled by randomness: The hidden role of chance in the markets and in life.* New York: Texere.

---

[i] Note that this is not the same thing as the more common stochastic model of *volatility*.

[ii] VaR – value-at-risk – is a risk measure now ubiquitous in banking and increasingly so in fund management as well. If the one-day, 95%, VaR for a portfolio is, say, -2.2%, that means that losses of greater than -2.2% should not occur more often than 1 day in every 20.