Off The Wall is our unpredictable monthly column of general interest articles, written by and for the scientific and financial communities. Here you can expect insightful and entertaining writing with an attitude. Sometimes technical, sometimes opinionated; but never dull. Check back often to see what's Off The Wall...and on our minds.

The market volatility during 3Q 1998 has resulted in spectacular and well-publicized losses, exemplified by the bailout of Long-Term Capital Management, and called into question the sophisticated trading models used in today's markets. Whether or not these models are to blame for these events, as the popular press would have us believe, there is little doubt that the Black-Scholes-Merton assumptions and their extensions are incapable of modeling large but rare fluctuations, or even the milder but statistically significant phenomena giving rise to the volatility smile and fat tails.

This month's Off The Wall, contributed by Erik Aurell, Roberto Baviera, Maurizio Serva, Ola Hammarlid and Angelo Vulpiani, reports on their recent work which attempts to use large deviation theory to model rare fluctuations, while at the same time reducing to standard option pricing theory for small fluctuations. We know that these issues are of great interest to the readership of Off The Wall, and extend a warm welcome to our latest contributors.

Large bets, rare fluctuations and derivative pricing

Introduction

The following text was essentially written up before the latest events on Wall Street involving the hedge fund Long-Term Capital Management. Since the words long-term capital have therefore been widely in the news lately, it was found appropriate to insert this preamble before the main body of the document.

The starting-point of our investigation was the scattered observations in the literature that speculative prices are not necessarily well modelled by traditional models in financial economics, i.e. do not necessarily have to be lognormal or other diffusion or diffusion-like processes. The theory of stochastic processes in probability theory (a branch of mathematics) contains an enormous collection of processes. Very few are similar to the ones widely studied and used in finance.

What if one of the non-standard ones was a better model? Typically, the market in, say, one such security and a bond, would be incomplete. The no-arbitrage arguments would then not suffice to price a derivative. This does not mean that arbitrage would be permitted in such market models. It only means, that if the market is not complete there are many different possible price systems, each of them without arbitrage. To actually know and compute the price of a derivative, the condition of no-arbitrage is still necessary, but it is not enough.

The following sections report on our attempts to find a viable approach to derivative pricing in incomplete markets. Detailed technical papers are in press, but we present here to a wider audience the salient ideas and conclusions. One conclusion is that standard financial models arise in a natural way as approximations. These approximations are valid for the most probable events in a very general setting, but are typically not valid for large and rare events. In concrete terms, the Black-Scholes formula is a valid approximation for most at-the-money options close to expiry in practically any model. But it can be wrong for in-the-money and out-of-the-money options. Indeed, such discrepancies are often observed, and referred to as "volatility smile" effects. In our theory, we can account for such behaviour in a comprehensive and concise manner, making few assumptions on the way.

The actual causes of the LTCM events are most likely complex, and will only become clear to the public over the coming months. However, if preliminary accounts in the general press can be taken at face value (*New York Times*, September 24, 1998 and September 26, 1998), then Mr. Meriwether and partners had placed bets that the spread between T-bonds and other bonds would not increase. We do not claim to have any insight on the strategies actually used at LTCM. But if this account is true, they had effectively bet that a rare event would not happen.

It is precisely in derivatives associated with rare, large and potentiallly catastrophic events that the methodology based on lognormality, Black-Scholes and totally hedgeable risk (even in principle), that this general approach is questionable.

Non-mathematical description and background

Our work proposes a theory of derivatives in incomplete markets, with or without transaction costs. Such a theory must proceed from some assumptions. A traditional set of assumptions in financial economics is that different economic agents are characterised by different attitudes to risk. The main draw-back of this idea is that it is not restrictive enough: without knowing the attitudes of agents actually in the market, such a theory has little predictive power. In derivative pricing theory, the main stress has therefore been on complete markets, where risk-return preferences do not matter. Real markets are not complete. Standard theories should therefore be considered as successful approximations, which could perhaps be systematically improved upon.

We postulate that agents in incomplete markets choose to optimize the long-term growth rate of their return on capital. The "bets" are investment strategies in underlying and derivative securities. As a concrete example, we assume that it happens many times that an operator wants to price an option, which is at-the-money and one month to expiration, and that the probability distribution of the underlying stays the same. Note that the operator may certainly bet on other gambles simultaneously. A kind of law of large numbers here applies, and it can be shown that there is one investment strategy which yields almost surely the fastest growth rate. After long times the return on capital using this strategy is almost surely larger than any other by an arbitrary amount.

The main limitation of the method should be clear from the preceeding paragraph: It must really be a good approximation that the same game is played many times. Otherwise we cannot use the law of large numbers, and we have to fall back on judging the merits of different investment strategies by risk-return arguments. Let us say that ten is a sufficiently large number, and assume that the probability distribution of the underlying is approximately stationary over the time-scale of months. The theory presented here then applies to an operator that faces the same investment choice several times a month. By "the same" we understand something very narrow, i.e.,.as in the example above, to price an option, which is at-the-money and one month to maturity, every time on an underlying governed by the same price process. In other words, this is a theory which is perhaps applicable to the rational choice of a full-time professional trader in an active market, but not to a private investor who rebalances his portfolio only occasionly.

The main advantage of the method is that it treats complete and incomplete market models in a unified manner. If the best and simplest fit to historical price data is an incomplete market model, then one can go ahead and use that model. A second advantage is, that, while it is certainly a technical problem, there is no conceptual problem to include transaction costs. As a bonus, leaving the realm of Ito calculus and Gaussian processes suggests new analytical tools from probability theory. Primarily we here intend the Large Deviation Theory (see e.g. Varadhan (1984)), which may have practical applications to pricing options far-out-of-the-money or well-in-the-money. Finally, one may derive the Black-Scholes model or other models using Ito calculus as approximations. From our point of view the Black-Scholes formula is so successful, not because prices of financial assets are actually lognormally distributed, which they may or may not be, but because the formula arises as a first approximation in a much wider class of models.

The idea that optimizing long-term growth-rates is a useful criterion for investments has a long intellectual history, almost all predating by many years the birth of modern financial markets. In fact, the basic idea was stated more than two centuries ago by Bernoulli (1738), and developped quite explicitly by Kelly (1956) in the early 1950's. It can safely be said that in those days, neither the data nor the computer equipment to effectively carry out the necessary optimizations was available. It was also probably not a very accurate assumption that a typical agent played the same game many times. A systematic critique of the growth-optimal criterion from the view-point of utility theory can be found in the well-known papers of Samuelson (1971) and Merton and Samuelson (1974). For more recent reviews the interested reader may look at Merton (1990)[expanding on his earlier views], or at Hakanson and Ziemba (1995) [defending the growth-optimal criterion]. Our input to this debate is that it boils down to a quantitative question whether the game really is played sufficiently many times, and the growth-optimal criterion is applicable. We further remark that while in simple models optimizing long-term growth rates is equivalent to maximizing expected logarithmic utility, in more realistic models this is not so. Logarithmic utility and long-term growth rates are in general different things, and arguments against one are not necessarily arguments against the other.

In the rest of this article we will give the salient points of our method. We will concentrate on the derivative pricing problem in an incomplete market. We will not touch here upon the problems of hedging, the inclusion of transaction costs, practical use of Large Deviation Theory or any but the simplest models. For discussions of more realistic models, as well as the other topics, technical details. equations and formulae, we refer to Aurell et al (1998a) and Aurell et al (1998b)

The Black-Scholes World

The Black-Scholes model and its discrete cousin, the Cox-Ross-Rubinstein model, assume certain forms of the probability distribution of a risky security. For the discussion that follows, the important point is that in both cases a market consisting of the risky security and one risk-less security (bond) is complete. Whatever happens in the future, say event X, one can construct a portfolio out of the security and the bond, which will be worth 1 if X occurs, and otherwise 0. When this is possible, portfolios in the security and the bond form a complete basis of the space of future states X.

A derivative security is something the value of which will be f(X) if X happens (to the security and the bond). If the market is complete we can put together a new portfolio in the security and the bond, which will be worth f(X), for every X. The familar argument of no arbitrage then fixes the price of the derivative today.

To smooth the transition to the next section and the world of incomplete markets, we will now write out the no-arbitrage argument in a way convenient to us. Assume that the prices of the derivative and the portfolio are different, and assume that an operator can invest a sum S without moving the market. Furthermore, assume that the operator would have to invest a sum M to move the market, and that S is much less than M. S would be less than M because of finite but large market depth, and of externalities for this operator.

Then the operator would bet S on the arbitrage opportunity, and in return he would get $(1+D)S$, where D is the price difference between the derivative and the portfolio. When the opportunity arises again he would do it anew, and after N such operations he would have $(1+D)^N S$. Very soon his capital would be larger than M. Then he would start to move the market, and the arbitrage opportunity would disappear.

A world with non-hedgeable uncertainty

In an incomplete market there can sometimes be pure arbitrage opportunities, but typically there are not. The return of investing capital S on a market game will be $(1+D(l))S$, where $D(l)$ is a random variable, depending parametrically on the strategy l, which we take to be a portfolio in the security, the derivative and the bond. If the operator does this many times, his excess returns in each consequtive bet will be, say, $D_1(l), D_2(l), D_3(l),\ldots$ and his aggregated return will be $(1+D_1(l))(1+D_2(l))(1+D_3(l))\cdots S$.

For this simple model we can take the logarithm, use the law of large numbers, and see that the capital grows in N bets to something which about as large as $\mathbf{exp}(NR(l))S$, where $R(l)$ is the long-term rate of return on capital using strategy l. A strategy that maximizes the growth rate will for large N almost surely yield a higher return than any other. Investors using this strategy will hence see their capital increase faster (almost surely) than others. If this strategy contains a positive amount of the derivative, the demand for the derivative will therefore grow. Hence the price of the derivative will tend to move so that the growth-maximization strategy contains no derivative.

In technical terms, we have derived a derivative pricing principle in incomplete markets by a particular equivalent Martingale measure. This measure obeys a variational principle, which roughly says that it is as similar as possible to the raw probability of the underlying, while satisfying the constraint that it is a Martingale.

Black-Scholes as an approximation

In the preceeding section we found the price of a derivative when betting on a game with uncertain outcomes. However, the set of possible outcomes was known in advance, like the returns in a lottery. This hence fixes e.g. the price of an option one elementary time-step before expiry. Having obtained this, we can work backwards in time, just as in the text-book discussions of the Cox-Ross-Rubinstein model, and deduce the price of derivative at an arbitrary earlier time.

The price of an option will now be the expected value upon expiry, the expectation value taken with respect to a particular Martingale measure Q. This measure has been obtained by aggregation under multiplication of elementary measures over each elementary time-step separately, each of which obeys a variational principle.

For large times, or, more precisely, for times T much larger than the elementary time-step, this measure Q must then be of the form that $Q(\log(S_T/S_0)/T)$ is about as large as $\mathbf{exp}(-TG(\log(S_T/S_0)/T))$. Expressions like these are probably familiar to many readers of this page as the ansatz of an intensive entropy function in statistical mechanics, or from the multifractal formalism in non-linear dynamical systems. In the present context the function G is called the Cramer function, and the behaviour exponential in T is called a Large Deviation formula. The meaning is, that under aggregation we arrive at a probability distribution over possible returns, $\log(S_T/S_0)/T$, and that this probability distribution depends in a simple way on T. We only need to determine the function G once, and then we are done for all sufficiently large T.

The minimum of the function G (maximum of the meaure Q) is obtained at some point x. If we expand around the minimum we have an approximate formula for Q as $\mathbf{exp}(-{}^1/_2\, cT(\log(S_T/S_0)/T - x)^2)$, where c is the second derivative of G at x. Since Q is a Martingale, the expected value of S_T under Q must be $\mathbf{exp}(rT)$, where r is the riskless return. We can use this constraint to solve for x in terms of c and r, and by insertion and explicit comparison, the resulting formula can easily be seen to be equivalent to Black-Scholes.

The moral of this derivation is that if the quadratic approximation is only good close to the minimum, it is incorrect to use that approximation to evaluate the Martingale constraint. It will also be incorrect to use that approximation to evaluate derivatives with linear pay-offs, i.e. options. In both cases, the reason is that the expectation value is dominated by the distribution of large and rare events, which are, by hypothesis, not distributed lognormally.

Conclusions

The growth-optimal criterion can be used as a simple and versatile tool to price derivative, making few assumptions on the probability distribution of the underlying securities, except what is directly available as empirical data.

With this method, derivative prices can be written as expectation values with respect to auxiliary risk-neutral probability distributions, just as in standard theory. The form of these auxiliary probabilities may however be very different from standard models. They do not have to be lognormal, or processes amenable to Ito calculus, or satisfy any simple PDEs.

All these generalizations point towards one simple fact: these distributions may have much massive support for large and rare events, which matter crucially one by one, and not averaged out in an unordered melee. When considering individual derivatives, one can use this theory to price far-out-of-the-money and well-in-the-money options, for which the lognormal distribution and the Black-Scholes formula are generally considered less accurate. It can also be used on portfolios of derivatives.

To end on a speculative note, a hedge fund with a mix of assets and derivatives chosen according to standard finance theory could be secured against moves that typically occur in the market, but not against very improbable events. This is only common sense. What our theory can do is to provide a frame-work in which to discuss such scenarios in quantitative terms. We hope that if the theory of derivative pricing in incomplete markets presented here will be taken seriously, it would also be of use in evaluation and containment of financial risk.

Acknowledgements

We thank our several collegues and collaborators in various projects that led to this work, and in particular Jean-Philippe Bouchaud for many enlightening discussions.

This paper is based on "Gambling and Pricing of Derivatives", published in the Financial Economics Network Working Paper Series, , and
"A general methodology to price and hedge derivatives in incomplete markets" , submitted to the International Journal of Theoretical and Applied Finance .

References:

1. E. Aurell *et al (1998a),* "Gambling and Pricing of Derivatives*, Financial Economics Network Working Paper Series*, May 26, 1998.
2. E. Aurell *et al (1998b), "*A general methodology to price and hedge derivatives in incomplete markets" , *International Journal of Theoretical and Applied Finance* [submitted],
3. D. Bernoulli (1738:1954), "Exposition of a new theory on the measurement of risk", *Econometrica **22***, 22-36.
4. R.C. Merton and P. Samuelson (1974), "Fallacy of the log-normal approximation to optimal portfolio decision over many periods", *J. Financial Economics **1***, 67-94.
5. R.C. Merton (1990), *Continuous-Time Finance*, Blackwell, Cambridge, MA.
6. J.L. Kelly, Jr. (1956), "A new interpretation of the information rate", *Bell System Technical Journal **35***, 917.
7. N. Hakanson and W. Ziemba (1995), *Capital Growth Theory*, vol. **9**, chap. **3**, 65-86, Elsevier Science.
8. P. Samuelson (1971), "The ``fallacy'' of maximizing the geometric mean in a long sequence of investing or gambling", *Proc. Nat. Acad. Sci. USA **68***, 2493-96.
9. S. Varadhan (1984), *Large Deviations and Applications*, Society for Industrial and Applied Mathematics, Philadelphia PA.

About the authors:

Erik Aurell is with Industrilogik L4i AB , Stockholm, Sweden, and with the Department of Mathematics, Stockholm University, Sweden. Roberto Baviera and Maurizio Serva are with the I.N.F.M. , Aquila, Italy, and Ola Hammarlid with the Department of Mathematical Statistics and Insurance Mathematics, Stockholm University, Sweden.. Angelo Vulpiani is with the I.N.F.M. , and with the Department of Physics ,University of Rome "La Sapienza", Italy.