Innovations in trading strategies

Izzy Nelken introduces some of the contributors to the pre-conference summit at ICBI Global Derivatives and Risk Management 2003.

Introduction

Last time I attended ICBI's Global Derivatives and Risk Management conference was in 1999. The mood then was quite different. Banks were gearing up to offer the latest 'exotic option' to their clients. There was a feeling that the good times would last forever.

Times have changed, however, and we are now suffering from a declining stock market and interest rates at extremely low levels. This has prompted many investors to look at alternative investments as an asset class that deserves much attention. In addition, the trading strategies themselves have become much more important.

What hasn't changed at all is the amount of talent and innovation that is present in this field.

With the recent developments in mind, ICBI has organized a one-day summit on 'Innovations in Trading Strategies', the theme being 'Quantitative, Relative Value and Proprietary Trading Strategies for Hedge Funds, Investment Banks and Asset Managers'[1].

Doyne Farmer looked at dynamic feedback between order placement and price formation. Armed with 350 million data points, he is building a model to explain bid–ask spreads and price diffusion rates. The model assumes that agents have zero intelligence and is reasonably accurate. The plan is to add intelligence to market agents on a piecemeal basis and observe the results.

Peter Carr considered derivatives on serial covariance. He first shows that the fair value for a serial covariance swap is always zero. It is quite trivial to hedge such an instrument. Rather than a swap, it would be useful to consider an option on the covariance. Such an option is difficult to hedge without a model. Thus the writer of such an option would be exposed to model risk. Therefore, a new hyper option is proposed. Every time the option is exercised the intrinsic value gets locked in and it changes from a call to a put and vice versa. This contract can be hedged by dynamic trading in European options. In addition, the price of a hyper option should be equal to a European. He further shows that European option prices need not be increasing functions of volatility.

Olivier Ledoit has done work on GARCH modelling; a particular sub-class called 'diagonal vech models'. These typically result in very large matrices. Even if one is willing to pay the computational price, the resulting covariance matrices may not be positive semi-definite. Ledoit uses decentralization to limit the size of the problem. The resulting estimators seem to perform better in practice.

Alireza Javaheri presented us with an overview of various filtering methods for volatility models. The main emphasis was on various types of Kalman filters. He applies these filters to various price data and shows how to best calibrate various models (e.g. stochastic volatility models). We can compare the implied volatility skew to the theoretical skew arising from the various models. Several trade ideas were discussed: selling puts and buying calls or selling far-term options and buying near-term options.

Lisa Borland has used a standard Brownian motion with statistical feedback to model stock price returns. Choosing a reasonable level of the fitness parameter resulted in an implied volatility option skew which is consistent with observations in the option markets, particularly for equity indices and currency options. The skews obtained from Borland's model are symmetric and do not correspond to implied volatility options on single stocks. The model will need to be modified somewhat to address this.

Richard Weissman showed several examples of mechanical trading systems. The structure, as well as the strengths and weaknesses, of such systems was explained.

Izzy Nelken concentrated on the liquidity premium associated with investments in hedge funds. How should investors judge a fund given that assets cannot be withdrawn for a long time period (often a year or two)? He first considered the payout profile for a hedge fund manager. This looks like a long position in a convertible bond. Next, he used numerical simulation to track two portfolios, one which is allowed to rebalance and the other which is not. The difference between these portfolios is the liquidity premium.

Matthew Dowle presented a pairs-trading strategy in stocks and the results of a historical back test. Much work has been done to choose the pairs and to test the viability of the strategy. On a historical basis, the strategy performed well in Europe and Japan over the test period 1992–2003, but failed to perform in the USA.

Antonio Mello has looked at how capital constraints affect traders' behaviours, in particular the interaction between an arbitrageur and other market participants that he deals with. These will typically be aware of the trader's position and also their risk capital constraints. When such counterparties sense that an arbitrageur is in a weak state they become like predators. They may lend money to the arbitrageur, allowing him to trade more and more and build larger and larger positions. This is in the hope of eventually forcing the arbitrageur to liquidate his entire position. It is then that liquidity will disappear and the predators will benefit from falling prices.

Tanya Beder gave a historical overview of quantitative trading. She gave some examples of such strategies. Her predictions for the future were quite interesting. She predicts the growth for certain styles of quantitative trading and the demise of other styles.

J-P Bouchaud spoke about volatility modelling. Earlier models assumed constant volatility. Later, stochastic volatility models, such as the Heston model, were developed. However, volatility seems to have complicated characteristics. For example, it appears that volatility fluctuations have long-term memories involving several time scales. Bouchaud is exploring the use of

[1] Many have presented joint work with others. In this introduction, the terms 'he' or 'she' will often mean 'they'.

the Bacry–Muzy–Delour model which may be able to capture some of the more salient features of the volatility process.

Richard Olsen told us that one of the main reasons there is a market is that different investors have differing investment horizons. Some wish to invest for a period of minutes while others buy and hold for many years. He and his group have analysed foreign exchange data and found differing levels of volatility that depended on the time of day.

Avoiding getting lost in the wilderness of bounded rationality: the path from zero intelligence to no arbitrage[2] by J Doyne Farmer

What really drives price formation? The standard neo-classical approach is to assume perfectly rational agents, and not worry much about the price setting mechanism. This approach has had striking successes in some areas, such as derivative pricing, but it has not been very successful in other areas, such as the pricing of underlying assets. In the last decade or so increasing attention has been given to agent models that perturb the perfect equilibrium model. Some of these generalizations have been useful. However, when it comes to the core problem of understanding the effect of bounded rationality, so far there has been little success in finding models with predictive power. The central problem is that there are an infinite number of boundedly rational models. People are complicated, and finding an accurate model of human behaviour is likely to remain elusive for some time to come. In addition, in many settings the equilibrium framework is sufficiently cumbersome to be intractable, and extending it is even harder. Many basic problems remain unsolved, such as a fundamental explanation for the variation in price volatility.

We take an opposite approach. We shed the heavy burden of modelling expectations in favour of a more realistic model of price setting. In particular, we focus our attention on the dynamic feedback between order placement and price formation. To be more specific, we model the continuous double auction under random-order placement. The key behavioural assumption is that agents respond to the current best prices and modify their orders accordingly. Prices respond to order placement, and order placement responds to prices. Some agents place market orders, and others place limit orders. The parameters of the model are the market-order rate, the limit-order rate density, the order-cancellation rate, the mean order size,

and the tick size. The model makes predictions about price volatility, bid–ask spreads, average stored liquidity in the limit-order book as a function of price level, time to fill, and any other statistical property of the market one can think of. To study the properties of the model we use computer simulation and analytic methods from statistical physics.

To test the model we use a unique data set from the electronic portion of the London Stock Exchange, which carries about 60% of the trading volume. This set contains roughly 350 million records, containing every action by every trader in every stock over a four-year period. From this data we are able to reconstruct the limit-order book for any given stock at any instant in time. Our analysis of this huge data set is just beginning, but we have already performed preliminary tests on about twenty stocks. We measure the five parameters of the model on a day-by-day basis and use these to make predictions about statistical properties, and compare the predictions to the measured sample values on the same day. The model explains roughly 60% of the variance in the bid–ask spread and price diffusion rate. We also study the average price impact of a market order as a function of its size, and show that in appropriately defined non-dimensional coordinates the price impact function for many stocks appears to roughly collapse onto a single curve. This result is particularly striking, as it suggests the existence of a universal excess demand curve.

It is very surprising to find this level of predictive power in a model in which the agents have zero intelligence. This research develops a line of thought originating with Becker, who suggested that the existence of a budget constraint alone dictates downward-sloping supply curves and upward-sloping demand curves. This idea was developed by Gode and Sunder, who showed that in the context of a trading experiment with pre-determined liquidation values budget constrained zero intelligence agents are reasonably efficient. Our model is in the spirit of earlier statistical models by Domowitz and Wang, and Bak, Paczuski and Shubik. We go beyond these earlier models by showing that, with appropriate modifications, they can make fairly accurate quantitative predictions about properties of real markets, predicting not just the sign of the slope but the quantitative details of supply and demand curves, transaction costs, volatility and liquidity.

We are in the process of extending these models by adding a little intelligence. One can view the framework of our model from an ecological point of view. The noise traders in our model can be thought of as liquidity demanders, willing to take small trading losses in order to get their trades done quickly. This creates profit-making opportunities for arbitrageurs, such as market markers. We believe that this can result in further improvements in the predictions of these models.

This inverts the usual approach in economics, by using zero intelligence random agents to probe market institutions and understand their consequences, and then adding a little intelligence. We thus work

[2] This discussion is based on the work of Marcus Daniels, Laszlo Gillemot, Giulia Iori, Fabrizio Lillo, Rosario Mantegna, Paolo Patelli, Eric Smith, Ilija Zovko, and myself.

our way incrementally from zero intelligence to no arbitrage. By starting with a zero intelligence model we avoid immediately getting lost in what Christopher Sims has called the 'wilderness of bounded rationality'. The zero-intelligence model tells us the best place to enter the wilderness. Once we are in it, it helps us keep from getting lost by providing a benchmark to measure what is due to market institutions and what is due to agent behaviour. The profit seeking nature of human utility tells us the right direction to walk. The final destination remains unknown.

References

Bak P, Paczuski M and Shubik M 1997 Price variations in a stock market with many agents *Physica* A **246** 430

Becker G 1962 Irrational behavior and economic theory *J. Political Economy* **70** 1–13

Daniels M, Farmer J D, Gillemot L, Iori G, and Smith D E 2003 Quantitative model of price diffusion and market friction based on trading as a mechanistic random process *Phys. Rev. Lett.* **90** 108102

Domowitz I and Wang J 1994 Auctions as algorithms *J. Econ. Dynamics Control* **18** 29

Gode D and Sunder S 1993 Allocative efficiency of markets with zero intelligence traders: markets as a partial substitute for individual rationality *J. Political Economy* **101** 119

Sims C 1980 Macroeconomics and reality *Econometrica* **69** 1–48

Smith E, Farmer J D, Gillemot L and Krishnamurthy S 2003 Statistical theory of the continuous double auction *Quant. Finance* submitted

After receiving a PhD in Physics, J Doyne Farmer worked in the theoretical division at the Center for Nonlinear Studies at Los Alamos National Laboratories, where he was Oppenheimer Fellow and founder of the Complex Systems Group. In 1991, he co-founded the Prediction Company, a firm whose business is automated trading of financial instruments grounded in time-series based directional forecasting methods.

Beyond variance swaps: trading auto-correlation[3] by Peter Carr

It is widely thought that path-dependent payoffs are highly sensitive to model risk. However, it is by now fairly well known that variance swaps can be robustly synthesized without specifying the process for volatility. In other joint work presented at the conference, we extend these results on variance swaps to contracts paying any function of the terminal price and realized variance. This extension allows for robust hedging and pricing of volatility swaps and options on variance for example. In the trading summit talk, we focus instead on the robust hedging of contracts whose payoffs depend on various measures of the realized auto-correlation of returns.

For example, we consider a serial covariance swap. Like a variance swap, this contract pays the difference between a floating payoff and a fixed payment, where the latter is chosen so that the swap is initially free. However, unlike a variance swap, the floating payoff is now the realized serial covariance of returns. When returns are specified as the daily percentage change in a futures price (with daily compounding), then the floating payoff can be replicated by trading futures daily. Unlike a variance swap, no options are required in the hedge. As a result, in a frictionless futures market, the fixed rate on a covariance swap is determined by no arbitrage and deterministic interest rates to be zero.

A potential drawback for investors con-

templating a position in a serial covariance swap is that the potential for losses is unlimited. This suggests introducing European options on realized serial covariance. As for any option, the maximum loss for a long position in a covariance option is limited to the initial premium. Long positions in calls on covariance would appeal to investors predicting trending, while long positions in puts on covariance would appeal to investors predicting reversion. Unfortunately, model-free replication of the payoff from European options on covariance appears to be tricky.

For this reason, we introduce a similar contract called a hyper option. A long position in a hyper option would appeal to investors predicting reversion. The reason is that a hyper option can be exercised multiple times; with each exercise converting a hyper call into a hyper put or vice versa. Since a hyper option can be exercised an unlimited number of times, all of the exercise proceeds are paid out in a lump sum at maturity (without earning interest in between exercise and expiry).

We focus on the valuation and hedging of a hyper option written on the forward price of its underlying asset. Hence, exercising a hyper call results in the difference between the forward price and the strike being paid to the owner at maturity. The owner also receives a hyper put which can be exercised at any time to lock in the difference between the strike price and the forward price. This difference is also received at maturity and the hyper put converts back into a hyper call.

A hyper option is clearly similar to an American option with the differences being multiple exercise and deferred receipt of the exercise proceeds. It is well known that a standard American option written on the forward price has a positive early exercise premium. The reason is that at any time prior to maturity, there is positive probability that the American option written on the forward price should be optimally exercised early. Similarly, we show that for a hyper option written on the forward price, there is positive probability that the hyper option should be optimally exercised early infinitely often. Despite this result, we also show that the hyper option has the same value as the corresponding European option. This surpris-

[3] Based on work co-authored with Roger Lee of Stanford University.

ing result is independent of the stochastic process governing the underlying forward price. The intuition behind the result is that every exercise policy is also optimal. Hence, one can value the hyper option under the simplest policy of waiting to maturity and exercising if the option is in the money then.

We present a model-free hedge for the writer of a hyper call. The hedge involves holding a European call when the hyper option is held as a call and holding a European put when the hyper option is held as a put. Hence, an exercise of the hyper option results in the hedger selling one European option and buying the European option of the opposite polarity. By put–call parity, the cash flow generated by this switch is just sufficient to buy bonds covering the liability locked in by the exercise. It follows that the fair premium to charge for the sale of a hyper option is just the initial premium of the corresponding European option required to initiate the hedge.

This equality in value between hyper and European options has implications for understanding the sources of value of a European option when markets are incomplete. For example, suppose that all investors have homogeneous beliefs and that they suddenly agree that (future) autocorrelation in the underlying forward price will be more negative. Then it is plausible that hyper option values would increase. By no arbitrage, European option values would also increase. This increase in option value occurs even though the statistical variance of the terminal log price decreases. Using a simple binomial example, we demonstrate that, indeed, European option values are not necessarily increasing in volatility, as has been demonstrated previously.

We conclude that there are at least a couple of ways to trade autocorrelation. Path-dependent contracts written on various measures of autocorrelation can be hedged and priced without a model for the dynamics of the underlying. Further research is needed on this ‘trendy’ topic.

References

Carr P and Lee R 2003 Hyper options *NYU Working Paper*

Peter Carr recently joined Bloomberg to head up a new group conducting quantitative research. He is also the new director of the Masters in Math Finance program at NYU’s Courant Institute. Prior to his current positions, he headed equity derivative research groups for six years at Banc of America Securities and at Morgan Stanley. His pervious academic positions include four years as an adjunct professor at Columbia University and eight years as a finance professor at Cornell University. Since receiving his PhD in Finance from UCLA in 1989, he has published extensively in both academic and industry-oriented journals. He is currently the treasurer of the Bachelier Finance Society and a practitioner director for the Financial Management Association. He is also an associate editor for six academic journals related to mathematical finance and derivatives. He has given numerous talks at both practitioner and academic conferences and was recently selected as Risk Magazine’s prestigious ‘Quant of the Year’ for 2003.

Flexible multivariate GARCH modelling[4] by Oliver Ledoit

Univariate GARCH(1,1) is the industry benchmark for modelling and forecasting time-varying volatility of the return on a single asset. If several assets are under consideration, a multivariate model is needed. Time-varying volatility is now measured by the conditional covariance matrix of the asset returns. Surprisingly, the extension of the GARCH(1,1) model to higher dimensions is a difficult undertaking.

The most natural multivariate model uses a separate one-dimensional GARCH(1,1) process for each conditional variance (of the return on any given asset) and for each conditional covariance (between the returns on any given pair of assets). This model is known as the *diagonal vech model*. It is intuitive, reasonable, and easy to understand. However, it is not commonly used. The main reason is a ‘curse of dimensionality’. Say we consider a universe of N assets. The diagonal vech model then requires N one-dimensional GARCH processes for the distinct conditional variances and $(N–1)N/2$ one-dimensional GARCH processes for the distinct conditional covariances. Each

process has three free parameters, so the total number of parameters to be estimated, $3N(N+1)/2$ increases rapidly with N. Unless $N \le 5$, multivariate optimization routines generally fail to converge. But there is a second reason. If N is small and the optimizer converges, one would then use the estimated model to forecast the conditional covariance matrix. Unfortunately, the underlying mathematics do not guarantee that this matrix will be positive semi-definite, that is, a ‘proper’ covariance matrix.

The common solution is to specify a restricted model. First, such a model will have fewer parameters and be easier to estimate. Second, the restrictions can be designed in such a way as to ensure the positive semi-definiteness of forecasted conditional covariance matrices. The feasibility of a restricted model comes at the expense of a high risk of misspecification. For example, the constant conditional correlation model postulates that the conditional correlation between the returns on any given pair of assets will not change over time. This assumption is clearly violated in equity markets where conditional correlations tend to increase during bear markets. Other restricted models produce other violations.

The innovation of our method is to go back to the general diagonal vech model and find a way to estimate it, even for $N > 5$. We do this by a *decentralization* of the estimation process. Instead of attacking the full model, we break the large problem into smaller pieces, each of which can be estimated easily. Next, we put the pieces together and obtain a joint estimation of the full model. But, still, forecasted covariance

[4] Based on work co-authored by Pedro Santa-Clara and Michael Wolf.

matrices may not be positive semi-definite. We solve this second problem with another novel step. It turns out that by transforming the matrices of estimated GARCH coefficients into positive semi-definite matrices, we can guarantee that the forecasted covariance matrices also will be positive semi-definite. Part of our contribution is to provide a numerical algorithm for this transformation step.

Using historical return data on $N = 7$ country stock indices, we compare our estimator to a number of existing ones. Using several comparison criteria (such as forecast criteria, testing for leftover correlation in standardized residuals, value-at-risk and portfolio selection), we find that the new estimator indeed yields improved performance.

References

Ledoit O, Santa-Clara P and Wolf M 2003 Flexible multivariate GARCH modelling with an application to international stock markets *Rev. Economics Statistics* **85** at press

After receiving his PhD from MIT, Olivier Ledoit served as an Assistant Professor of Finance at UCLA. He has been awarded the Roger F Murray prize by the Q Group for his applied work on covariance matrices. His research spans the areas of finance, statistics and economics. He is currently Head of the Statistical Arbitrage Desk in the Global Equity Proprietary Trading Group of Credit Suisse First Boston.

Identifying statistical arbitrage opportunities and enhancing volatility trading using a new approach to time series analysis[5] by Alireza Javaheri

In my talk I presented an introduction to various filtering algorithms: the Kalman filter, the extended Kalman filter (EKF), as well as the unscented Kalman filter (UKF), similar to Kushner's nonlinear filter. I tackle the subject of non-Gaussian filters and describe the particle-filtering (PF) algorithm. In particular, the extended particle filter (EPF) and the unscented particle filter (UPF) are discussed in detail. In addition to the background and theory, a step-by-step algorithm is provided for each filter and I also explain how to find the optimal model parameters via the maximization of the likelihood (ML) function, both in the Gaussian and non-Gaussian case.

I then apply the filters to a term-structure model of commodity prices and the main results are as follows. Firstly, the approximation introduced in the extended filter has an influence on the model performance. Secondly, the estimation results are sensitive to the system matrix containing the errors of the measurement equation. Thirdly, the approximation made in the extended filter is not a real issue until the model becomes highly nonlinear. In that case, other nonlinear filters such as those described in the first section of the article may be used. The performance measurements are carried out via the mean pricing error (MPE) and the root-mean square error (RMSE) valuations.

Lastly, the application of the filters to stochastic volatility models shows that the particle filters perform better than the Gaussian ones; however they are also more expensive. What is more, given the tests carried out on the S&P 500 data it seems that, despite its vast popularity, the Heston model does not perform as well as the 3/2 representation. The robustness of each model is also considered. This suggests further directions for research on other existing models such as jump diffusion, variance gamma or CGMY. Clearly, because of the non-Gaussianity of these models, the particle filtering technique will need to be applied to them.

Another idea would be to compare the implied parameters from the options in a cross-sectional manner at a given point in time, with those obtained from filtering the above time-series during the life of those options. In order for the parameters to be comparable the Girsanov theorem is applied to the time series and everything is considered in a risk-neutral framework. A difference between the parameter sets could indicate either a misspecification in the model, or a trading opportunity. Others such as Yacine Aït-Sahalia have already looked at the difference between cross-sectional and time-series implied distributions; however they have used a non-parametric methodology without the assumption of stochastic volatility. The advantage of the parametric approach is that we are able to distinguish between volatility-drift parameters on the one hand and volatility-of-volatility parameters on the other. However, it does have to make an assumption on the form of the volatility process.

References

Aït-Sahalia Y, Wang Y and Yared F 2001 Do option markets correctly price the probabilities of movement of the underlying asset? *J. Econometrics* **102** 67–110

Javaheri A, Lautier D and Galli A 2003 Filtering in finance *Wilmott Magazine* (5)

Prior to joining RBC Capital Markets as a Quantitative Analyst in Global equity derivatives, Alireza Javaheri held positions at Goldman Sachs and Lehman

[5] Based on the article co-authored by Delphine Lautier and Alain Galli.

Brothers. He holds an MSc in Electrical Engineering from MIT and is also a CFA Charter holder. He is currently working on his PhD Thesis on stochastic volatility at Ecole des Mines de Paris with Professor Alain Galli.

Non-Gaussian option pricing by Lisa Borland

Options are important financial instruments traded in huge volumes on global exchanges around the world. They are clearly important for hedging purposes, potentially protecting investors from large losses due to uncertain events. Information from the options markets can also be a leading indicator to investors with trading strategies based on the underlying stock or commodity. Correctly spotting mispricings in options themselves can also lead to potential gains. Being able to price options correctly and the ability to extract meaningful information from market prices is therefore clearly a first step to exploiting the opportunities that the options markets offer.

Since its introduction in the early 1970s, the Black–Scholes (BS) option pricing equations have been the basis for traders to theoretically obtain fair values of options and other derivatives. Nevertheless, it is well known that the BS prices typically deviate from empirically traded prices in a rather systematic way. In particular, to be able to reproduce market prices, a different value of the volatility of the underlying stock must be used for each value of the strike price of the option. A plot of implied volatility versus strike usually forms a convex function, dubbed 'the volatility smile'. Traders are very used to working with volatility smiles and the notion of a smile is widely accepted. However, it must be an artefact of false assumptions in the BS approach. Indeed, one would expect the volatility to have only one value, perhaps varying depending on the relevant time horizon of the option, but certainly not on how much profit is potentially to be made in a particular transaction.

One of the main assumptions of the Black–Scholes theory is that stock returns are log-normally distributed. While this greatly simplifies the mathematical treatment of the problem, it is widely recognized that real returns exhibit fat power-law tails. This model misspecification is certainly one source of the discrepancy between theoretical option prices and market prices. In the present work we go beyond the log-normal world of Black-Scholes, by introducing a new returns model with non-Gaussian driving noise. Fluctuations are assumed to evolve anomalously in time according to a non-linear Fökker–Planck equation. Such a noise process can be modelled as the result of a standard Gaussian process with statistical feedback. The degree of feedback is characterized by the index q, such that if $q > 1$, then rare events will lead to large fluctuations, whereas more common events will result in more moderate fluctuations. This model of returns is consistent with empirical observations of returns distributions. For example, $q = 1.4$ results in a fat-tailed non-Gaussian distribution which fits well to the distribution of individual stock returns.

Many of the nice mathematical properties of the standard Gaussian noise process are preserved for $q > 1$, and it is possible to derive analytical formulae for finding the fair value of options based on this model. Closed form solutions are obtained for arbitrary $q > 1$, recovering the usual BS if $q = 1$.Theoretical option prices predicted by this model using $q = 1.4$ are in good agreement with empirically observed option prices. In particular, using just one value of the volatility across strikes, we are able to reproduce market prices quite closely. Instead of a curved volatility surface which is needed to price options across strikes and times to expiration, we need just one number across strikes. This has important implications, especially for volatility trading strategies.

References

Borland L 2002 A theory of non-Gaussian pricing *Quant. Finance* **2** 415–31

Dr Lisa Borland is a senior researcher at Evnine-Vaughan Associates, a San Francisco-based hedge fund operating a quantitative statistical arbitrage strategy. Her work is focused mainly on equity and options markets, with an emphasis on developing new trading strategies.

Developing successful mechanical trading systems: a brief synopsis by Richard L Weissman

A definition of terms

Mechanical trading systems can be defined as a method of generating trading signals and quantifying risk that is independent of an individual trader's discretion. The type of mechanical trading systems that I have developed and traded over the years are all based on mathematical technical analysis. Mathematical technical analysis can be defined as the utilization of past price history in order to forecast future price trends and identify low-risk/ high-reward trading opportunities.

Trend-following and counter-trend trading systems

Most (around 70%) of the time, prices trade in a sideways or 'range-bound' pattern. By contrast, markets are only in a 'trending' mode around 30% of the time[6]. In statistical terms, commodity and financial markets

[6] Based on proprietary studies I have found that the majority of trading instruments are range-bound roughly 70% of the time. Of course certain markets such as foreign exchange cross rates display an even greater propensity towards mean reversion, while others like short-term debt instruments show greater tendencies to trend.

are said to be leptokurtic. That is, they display a strong tendency towards mean reversion[7]; in other words, prices tend to cluster around the mean.

Yet a large portion of mechanical trading systems are dedicated to trend identification because when prices are not in this mean reversion mode, they tend to trend[8]. In statistical terms, commodity and financial markets are said to be leptokurtic with fat tails, meaning that when they are not in their mean-reverting mode, they tend to display powerful and sustainable trends. These trends offer traders low-risk/ high-reward opportunities, such that a single profitable trend-following trade will often offset numerous small losses, thereby resulting in an overall profitable trading system that experiences win/loss ratios of less than 1:1.

Improving the rate of return

There are two methods commonly employed to improve overall rates of return without significantly increasing risk: the addition of various non-correlated assets, such as crude oil and foreign exchange futures, into a single trading system and the combination of

non-correlated trading systems[9] within a single fund.

Benefits and limitations

One of the most obvious benefits in using mechanical trading systems is that the results generated are objective and indisputable. As a result traders can determine both the potential risk/reward ratio and probability of a particular trade's success, based on a study of the system's historical performance, prior to any commitment of capital.

The most serious limitation in utilizing mechanical trading systems is the weathering of losses that experienced discretionary traders can often avoid. For example, following the resolution of the 2003 conflict in the Middle East, the vast majority of trend-following technical indicators generated buy signals in the US Dollar. However, many traders disregarded these signals because they represented market emotion and were in conflict with a weakened supply-and-demand picture.

Conclusion

Within the context of these limitations, mechanical trading systems remain an invaluable tool. Since trading systems generate buy and sell signals based upon a different criteria than discretionary traders, their results should have a low correlation with those of discretionary trading books. If both trading methodologies are profitable overall, the incorporation of mechanical trading systems into a discretionary trading environment should smooth out worst peak-to-valley drawdowns in equity, reduce durations of equity drawdowns and produce overall improvements in return on investments.

References

Murphy J J 1999 *Technical Analysis of the Financial Markets* (Englewood Cliffs, NJ: Prentice-Hall)

Richard Weissman is a trading consultant and faculty member of the Oxford–Princeton Programme in Princeton, NJ, USA. He is based in Florida. E-mail: rweissman@oxfordprinceton.com.

[7] Oscillators such as Wilder's Relative Strength Index or Bollinger Bands attempt to profit from this tendency towards mean reversion. For explanations of various oscillator formulas see Murphy (1999).

[8] Various moving average indicators form the building blocks for many trend following trading systems.

[9] Mean reversion trading systems and trend-following systems.

A liquidity haircut for hedge funds[10] by Izzy Nelken

One of the most important issues an investor faces is that of building a portfolio of different assets. What percentage of the portfolio should be allocated to each asset? Since the 1950s mean variance optimization has been used to construct portfolios of traditional assets such as stocks and bonds. As alternative investments grow in size and importance, more investors add this asset class to their portfolios. It is natural to ask how much should be allocated to a specific hedge fund or to a fund of funds.

Many hedge funds have outperformed both bonds and equities on a risk-adjusted basis. Therefore, when one naively applies the mean variance framework to hedge fund investments, one gets an over-allocation to

alternative investments. Many practitioners simply cap the allocation to hedge funds at some 'reasonable' level which is determined *ad hoc*. This reflects the ideas that hedge funds have inherent risks that are not captured in a mean variance framework.

One of the main risks is that of liquidity. An investor in a hedge fund is usually prohibited from withdrawing the funds until the expiration of a lockup period which is typically a year or two. Our work 'penalizes' the historical volatility of a hedge fund to account for the non-liquidity.

The two main steps are as follows

Consider the payoff to a hedge fund manager. Previous papers claimed that a hedge-fund manager is long a call option on the fund. We modify this to say that a hedge-fund manager is long a convertible bond. If the fund appreciates, the manager gets richly rewarded. However, if the fund depreciates

[10] Based on work co-authored by Hari Krishnan, vice-president in the investment management division of Morgan Stanley.

beyond a certain level, then investors may withdraw capital from the fund. This will cause the fund manager to lose not only this year's revenues but revenues from future years. This is similar to an investor who owns a convertible bond whose issuer has defaulted. The fund manager adjusts their leverage so as to maximize the expected return from the asset.

We then compare two cases. In each case a portfolio is constructed consisting of a reference asset (such as a money market account) and a hedge fund. The first case has unrestricted liquidity between the two assets and the investor can constantly rebalance. In the second case, we assume that no rebalancing is allowed. The investor is always trying to maximize a utility function such as 'portfolio return/portfolio volatility'.

As the hedge-fund manger modifies the leverage in a way that depends on the underlying asset price, we cannot get closed-form solutions. Instead, we rely on numerical simulation. For a particular example case, we obtained that the historical volatility of the hedge fund should be amplified by 10% to account for the liquidity premium. Thus, if the historical volatility was 20%, the input to the mean variance optimizer should be 22%.

Our approach has been to modify the existing mean variance optimization approach to allow for inclusion of hedge funds. It is unclear whether this approach should be applied to alternative investments at all. The asset question as it pertains to hedge funds should be studied further.

References

Nelken I and Krishnan H 2003 Haircutting the hedge funds *Risk Magazine* April

Izzy Nelken is President of Super Computer Consulting, Inc. at 3943 Bordeaux Drive, Northbrook, IL 60062, USA, www.supercc.com

Determining the profitability of pairs trading strategies[11] by Matthew Dowle

Pairs trading is a market neutral statistical arbitrage tool used by both hedge funds and long-only investors. It is based on the idea that two closely related securities will continue to be related in the future. If the relationship is temporarily violated then profit opportunities arise by buying one security and short-selling the other. Gains are realized when the price relationship is restored. This strategy has a long tradition in arbitrage, typically involving different classes of shares of the same company, e.g., ordinary/preference, ADR/underlying, etc. Our aim, however, is to identify different company pairs which may not be in the same industry or market and are not necessarily obvious. These are generally less sensitive to intra-day movements and realize higher returns than same-company pairs.

The 200 most liquid stocks in Europe form 19 900 exhaustive pairs. Identifying which pairs to trade is therefore a challenge. We choose pairs in the same sector but not necessarily the same industry, which have similar betas, and pass the Dickey–Fuller test for stationarity of the log price ratio at the 99% confidence level. We believe that mean reversion, not correlation, is the key to a successful strategy. As of March 2003, around 110 of the 19 900 potential pairs are chosen in Europe. We then open a position in a pair when the log ratio of the two prices diverges by more than two standard deviations, and close when the ratio reverts back to the rolling 130-day mean. A 20% stop loss and maximum 65-day holding period is imposed. An important additional risk control is waiting until after the log ratio has turned and crossed back through two standard deviations, towards the mean, before opening a position. This avoids opening long positions on distressed stocks for example. Model parameters were chosen arbitrarily and have remained unchanged since the original model specification in October 2001.

Back-testing the strategy is a non-trivial exercise. A different number of pairs pass the selection criteria every month providing a varying number of opportunities with varying durations. We test the strategy strictly out-of-sample from 1992 to 2003 in Europe, the US and Japan and report daily mark to market net asset value in USD of a pairs fund which follows all signals and aims to keep 80% invested in pairs, 20% cash. The cash balance reduces during runs of new signals and increases when positions are closed. No leverage is applied. Dividends are paid/received on short/long positions. We find average annual returns of 15.7% and 17.4% in Europe and Japan respectively with less than half the volatility of the local market. The maximum monthly loss over the period was –9.4% in Europe and –7.2% in Japan. The performance of the strategy is consistent and has very low correlation to the market. Annualized Sharpe ratios after 25 basis points cost per trade were 0.77 in Europe and 0.64 in Japan, using the T-Bill risk free rate. In the US however, we find that costs wipe out profits over the test period, resulting in a Sharpe ratio of –0.45. This corroborates observations by Goetzmann, Gatev and Rouwenhorst, who saw returns tailing off at the end of their 1962–1997 US pairs trading test. Further research is required as the methods are not directly comparable.

We have also applied the pairs trading strategy to manage an enhanced index fund with low tracking error. The strategy involves a core portfolio closely tracking MSCI Europe and a small part of the fund actively managed using pairs trading. Using 50 bps bet size and constrained weights we find that the enhanced index strategy outperforms the benchmark by 1.2% per annum with a tracking error of 70 bps. When we trade at the close of the day following the open or close signal, we find this only drops to 0.9%

[11] Joint work with Manolis Liodakis and Ciprian Marin.

out-performance with 68 bps tracking error. Therefore the strategy does not appear to be sensitive to achieving the closing price on the day of the signal.

In summary, we found significant risk adjusted returns after costs in Europe and Japan between 1992 and 2003, but not in the US. Further research on this regional difference would be worthwhile, not only for pairs trading but for gaining some insight into how similar hedge fund strategies may perform in future.

References

Liodakis M, Dowle M and Marin C 2002 Pairs trading for the risk averse (Smith Barney) March 2002

Liodakis M and Dowle M 2001 Pairs trading: a way to profit in a volatile equity market (Smith Barney) October 2001

Gatev E, Goetzmann W N and Rouwenhorst K G 1999 Pairs trading: performance of a relative value arbitrage rule *Yale School of Management Working Paper*

Matthew Dowle is Vice President in the Global Quantitative Research group of Citigroup Smith Barney. Much of his time has been spent developing the pairs trading model and servicing the firms' clients who execute this strategy. Other responsibilities include the firm's risk attribute model for European equities. The team was ranked first in the 2002 Global Institutional Investor survey for quantitative analysis.

Arbitraging arbitrageurs[12] by Antonio Mello

In many financial and commodity markets, trading is dominated by a few large and visible arbitrageurs. With their sophisticated analytical models, arbitrageurs exploit temporary price deviations, and by doing so supply liquidity much needed in order for the markets to function smoothly. The ability to quickly correct market imperfections rests on a solid capital base. Without it, large arbitrageurs become immobilized, unable to profit from price deviations and constrained from reducing exposures. Once seen as stabilizing forces, weak arbitrageurs can contribute to destabilize markets when regulatory constraints and margin requirements force them to liquidate their holdings over short periods of times. Arbitrageurs are then vulnerable to the actions of traders who conspire against the weak.

This talk discusses the strategies of traders who prey on financially weak arbitrageurs. Most studies on arbitrage trading ignore that some traders spotting a weak arbitrageur ride the market until they decide to pull the plug and precipitate a collapse. The fall of names such as Mettalgeselschaft, LTCM and AIB-Allfirst show how large arbitrageurs can fall victim to the actions of predators. Strategic trading increases the risks and deters arbitrage.

In the course of trading, arbitrageurs interact with many other traders. Brokers, lenders and counterparties are all well placed to learn about the arbitrageur's financial condition. When the arbitrageur is well capitalized, these traders just go about their normal business, but once they spot a fragile arbitrageur, they suddenly become aggressive. An example is when an arbitrageur, risking violating its minimum capital requirements, needs to regain control of the situation by liquidating positions. In less than perfectly liquid markets, this is costly, as liquidation trades have an adverse impact on price. More important, such trading becomes predictable. Strategic traders can then use their knowledge of the financial condition of the arbitrageur to benefit from the predictability of price deviations caused by the actions of the constrained arbitrageur.

At times when liquidity dries up, the adverse price impact caused by arbitrage liquidation can be so great that controlled liquidation of positions by an arbitrageur is no longer viable. Then, the arbitrageur must either liquidate everything, or use its size to move prices in an attempt to ensure its own survival. Price manipulation is, however, costly. It involves trading at an unfavourable price—an arbitrageur with a long position needs to buy more of the asset to keep the price high, while an arbitrageur with a short position needs to sell more to keep the price low. Again, deviations of price from value create an incentive for strategic traders to enter the market and take the other side of the arbitrageur's trade.

Strategic traders decide their trades based on the expectation of the arbitrageur's actions. When there is considerable uncertainty regarding what the arbitrageur will do next, whether he liquidates or buys to support prices, strategic traders may wait and not trade immediately, and instead lend to the arbitrageur. They lend even if the loan *per se* has a negative expected return, due to the low credit standing of the arbitrageur. The reason is that by extending credit to the arbitrageur, the strategic trader allows the arbitrageur to continue trading, and therefore gain from selling the asset at the artificially high prices that result as the arbitrageur with a long position buys more of the asset in a gamble to survive. And if the gamble for survival does not pay off, the strategic trader also benefits from a collapse in prices that results when the arbitrageur is forced to dump the resulting larger positions. Even if the arbitrageur succeeds in ensuring its survival, the strategic trader benefits from the lower prices that result as the arbitrageur trims positions back to normal levels. It is the action of providing finance to and simultaneously trading against the arbitrageur that makes the strategic trader behaviour predatory. The link between the plan of forcing a collapse and the importance of financing constraints for the success of such a trading strategy becomes clear in the predators' behaviour. The predator takes advantage of depressed prices during the time of the collapse, as well as the arbitrageur's price manipulation efforts before the collapse.

[12] Based on work co-authored by Mukarram Attari and Martin Ruckes.

References

Attari M, Mello A S and Ruckes M E 2002 Arbitraging arbitrageurs *University of Wisconsin-Madison Working Paper*

Antonio Mello is Professor of Finance at the University of Wisconsin, Madison. He is a former Managing Director of BCP Investimento, the largest bank in Portugal, a former Director of Finpro, a European private equity firm and sits on the investment committee of Norfin, a real-estate investment trust. He is past President of the European Financial Management Association and a Board Member of the Financial Management Association US. He is a research fellow of the Center of Economic Policy Research in London and was previously Chief Economist of the Central Bank of Portugal, as well as a member of the Subcommittee of Monetary Policy of European Union Central Bank Governors. He has previously taught at the MIT Sloan School of Management, the London Business School and INSEAD, France.

The past, present and future of quantitative trading by Tanya Beder

Quantitative trading has seen a dramatic increase in volume over the last few years despite the overall downturn in many financial markets. In this presentation, I cover three subject areas.

- Just what is quantitative trading?
- The past and present of quantitative trading.
- Predictions for the future.

Quantitative trading is a huge and growing area in both traditional and alternative investment programs. It is so broad that often ten people in a room will have eleven different definitions for what constitutes the business. Characteristics common to successful quantitative trading programmes are consistency (the same information always produces the same trading signals); discipline (signals are never 'over-ridden'); and persistence (no breaks or vacations!). In this presentation I review several examples of quantitative trading systems, ranging from those driven by purely fundamental factors to those driven by purely technical factors and combinations of the two. The almost infinite possibilities for the design and implementation of quantitative trading systems becomes apparent as a galaxy of styles are revealed and evaluated in six dimensions.

- Market (foreign exchange, fixed income, commodities, equities, other).
- Sector (spot/cash, forward, derivative, index, etc).
- Region (North America, European, Asia, G-7, Emerging, etc).
- Strategy (trend based, global macro, arbitrage, market neutral, long/short, etc).
- Math and model type (too numerous to list).
- Trade signal cycle (from low to high-frequency).

The 'past' of quantitative trading was a golden age that was not very long ago (mid 1940s to the mid 1970s) with many notable achievements. I review the key drivers of the evolution of the business in data collection and storage, mathematical models, financial methods, computation and trade execution. These are contrasted to the present business of quantitative trading and the dramatic advancements that have taken place. Specific examples are provided of current-day quantitative trading systems. The past and present vary substantially in terms of the dimensions of the business that experienced the greatest advancements.

In the final part of the presentation I lay the groundwork for why quantitative trading is likely to grow at an even faster pace in the future. Its low correlation to traditional investing as well as to other forms of alternative investing, its scalability, and its ability to perform through numerous market environments are some of the factors discussed. The areas of quantitative trading I predict will have the greatest growth include the following.

- High(er) frequency trading styles.
- More complex intra-day volatility strategies.
- More complex intra-day correlation strategies.
- Trend following and several forms of global macro.
- Arbitrage, especially in second- and third-order strategies.

Tanya Beder is a Managing Director of Caxton Associates, LLC, a multi-billion dollar investment management firm, where she heads the Strategic Quantitative Investments Division. Prior to Caxton, Ms Beder had over 20 years of Wall Street experience as President of Capital Market Advisors and as Vice-President of The First Boston Corporation. In 1997 Euromoney named Ms Beder as one of the top 50 women in finance around the world. Ms Beder is currently on the faculty of the Yale School of Management. She is also Chairman of the IAFE.

New models of volatility by J-P Bouchaud

Understanding and modelling the dynamical fluctuations of volatility in financial markets is of crucial importance both for risk control purposes and for option pricing and smile prediction. The statistics of volatility fluctuations is highly non-trivial since (a) the unconditional volatility distribution is found to be close to a log-normal (or an inverse gamma distribution), (b) volatility fluctuations have long-term memory, that cannot be characterized by a single time scale, but rather by at least three different time scales (days, months, years) and (c) there are causal, asymmetric correlations between past returns and future price changes that lead to skewness in return distributions [1]. Popular models of stochastic volatility, such as the Heston model [2], factor in these features in a qualitative manner, but fail to be quantitatively faithful to empirical data. Most importantly, point (b) above is not captured by the Heston model, within which all correlation functions decay exponentially with time. A new, promising path is to use the Bacry–Muzy–Delour multifractal model [3], inspired by turbulence data and also by early 'cascade' models by Mandelbrot. In this framework, the long-range, scale-free memory of the volatility process can be accounted for; variable range skewness effects can also be included [4]. The price process has multiscaling properties, which means that different moments of relative price changes all scale as a power-law of the time lag, but with different exponents.

In this way, both the skewness and kurtosis of the price distribution can be given non-trivial (power-law) dynamics as a function of maturity T with a tuneable exponent, and explain why the long maturity smiles are still found to be curved and skewed. This feature that cannot be captured by single-time scale models, where the skewness is found to decay as $T^{-1/2}$ and the kurtosis as T^{-1}, is much too fast to explain the curvature of long-term smiles. Finally, the distribution of drawdowns can be computed in the BMD model, and is found to be a power

law for large drawdowns. This decay is much slower than that found within the Heston model, say, and illustrates the importance of the volatility long term memory for the control of long-term risks. This result can be used to monitor the leverage of hedge-fund strategies (as is implemented in CFM) such as to minimize the probability of deep drawdowns.

References

[1] Bouchaud J-P and Potters M 2000 *Theory of Financial Risks* (Cambridge: Cambridge University Press)

[2] Heston S 1993 *Rev. Financial Studies* **6** 327–43

[3] Muzy J-F, Delour J and Bacry E 2000 *Eur. Phys. J. B* **17** 537–48

[4] Pochart B and Bouchaud J-P 2002 *Quant. Finance* **2** 303–14

Jean-Philippe Bouchaud was born in 1962 in Boulogne-Billancourt (France). After studying at the French Lycée of London, he graduated from the Ecole Normale Supérieure in Paris, where he also obtained his PhD. After a year spent as a research associate at the Cavendish Laboratory (Cambridge), he joined the Service de Physique de l'Etat Condense (CEA-Saclay), where he worked on the dynamics of glassy systems and on granular media. Together with J P Aguilar, he founded the company Science & Finance in 1994, now the research branch of Capital Fund Management that he supervises. Together with Marc Potters, he wrote the book Theory of Financial Risks in 1997. He was awarded the IBM young scientist prize in 1990 and the CNRS Silver Medal in 1996.

Meeting the challenge of developing volatility-forecasting models to characterize short- and long-term series behaviour by Richard Olsen

This talk addresses the understanding of the needs of the buy-side demand for intra-day options benchmarking the performance of option pricing system requirements cost–benefit analysis.

The development of pricing models of derivative instruments is driven by the demand of the buy side. In recent years, we have observed a growing interest in digital options and more recently in intra-day options. With digital options the payout is price path dependent and thus highly sensitive to the quality of the underlying volatility forecast. Intra-day options cannot be priced with volatility models using daily data. They require intra-day volatility models. Practitioners have been slow to introduce intra-day volatility models due to their complexity. There is a growing academic literature of the statistical properties of tick-

by-tick market data. This research plays an important role for intra-day option pricing models.

In the past, volatility-pricing models have been computed on the basis of daily prices. Traditionally, it has been argued that tick-by-tick market prices are noise which does not contain any relevant information. High-frequency finance tells a different story. Market prices exhibit a 24-hour seasonality, which is the result of the world turning on its own axis and market participants following their daily rhythm of going to work, having lunch and returning home. There are three distinct trading centres: Asia, Europe and America. In each of the three regions there are peaks of volatility reflecting the amount of trading that goes on in the particular financial instrument.

With daily data, the highest rate of innovation is 24 hours. If a market shock occurs within the 24-hour interval, it will take until the next price update for this information to be processed. The daily price update is just one observation and thus subject to a large error. With tick-by-tick data, it is possible to continuously incorporate new information. The information content is thus far richer than with daily data.

The continuously updated volatility measurements are, however, only meaningful, if the underlying 24-hour seasonality of volatility is accounted for. By processing tick-by-tick data and incorporating the '24-hour seasonality effect' of volatility it is possible to improve the quality of volatility models by a factor of three or more.

What are the costs of such a project? We estimate the cost to be in the range of 0.5 to 1.0 Mio USD. The project comprises the tasks to create a tick-by-tick data repository, develop an environment to compute the business time scale to account for the 24-hour seasonality, establish a testing environment and implement an online implementation of the pricing algorithm. The move to tick-by-tick data increases the complexity of the project. Minor issues, such as data quality, become major issues, because they have to be fully automated. With daily data, manual checking is feasible. With tick-by-tick data this is no longer possible. In planning the project, a cost–benefit analysis has to be done of how much of the technology has to be built in-house or bought from external providers.

How can such an investment pay off? First of all, there is an information advantage. Only a few market participants have started to use higher frequency data to estimate their volatility models. In this sense, the move to tick-by-tick data poses a competitive advantage in terms of 'informational edge' and thus generates added profit opportunity. Finally, certain products, such as intra-day options, cannot be offered without high-frequency data.

Richard Olsen is the Chief Executive of Olsen Ltd, a hedge fund manager specializing in currencies. He is also the founder of OANDA Corporation.

Now in its 10th successful year, ICBI's Global Derivatives event brought together over 350 senior industry figures from more than 25 countries to discuss the latest issues in quantitative finance and risk management. The unrivalled speaker faculty for the event comprised over 75 leading global financial minds providing unparalleled insights into the latest innovations in derivatives pricing, hedging and trading. For more information on how to get involved with Global Derivatives 2004 please contact Julian Kirby on +44 207 915 5072.

ICBI