GLOBAL
STRUCTURED CREDIT
RESEARCH

Structured Credit Products

20 AUGUST 2004

GLOBAL

David Shelton
(44-20) 7986-9607
david.shelton@citigroup.com
London

Back To Normal

Proxy Integration: A fast accurate method for CDO and CDO-Squared pricing

- **The rapid growth of CDO and CDO-Squared has created a need for fast sophisticated models**
- **We present a new method for quickly and accurately determining price and risk for both CDO *and* CDO-Squared trades**
- **In our tests with real trades and market data, we found the accuracy of Proxy Integration to be comparable to standard Monte Carlo, but with the practical advantage that computational speed was markedly better**

Contents

Executive Summary 3

Proxy Integration **4**

Overview 4

Valuation of CDO and CDO^2 in terms of the loss distribution 5

An accurate proxy for the joint loss distribution 5

Determining the expected loss on a tranche using the Proxy Distribution 8

Results 9

Conclusions 12

Executive Summary

Rapid growth in synthetic CDOs

One of the biggest recent developments in credit derivatives has been the growth of the single tranche synthetic CDO market. Standardised CDX/ITRAXX tranches have also added liquidity and transparency to this market.

Evolution of CDO-Squared structures

The search for yield in a tight-spread environment has also fuelled interest in more complex leveraged structures, such as CDO of CDOs ('CDO-Squared' or CDO^2). These new developments have created a need for sophisticated models to generate mark to market and risk in a timely and accurate fashion.

Existing numerical techniques are slow

The Gaussian copula has become the market standard model for quoting these structures. Traditionally, because of the large numbers of underlying credits involved, these models have been implemented in Monte Carlo. This method has several disadvantages: the risks are inaccurate, and it is computationally expensive.

Fast numerical implementation for CDO pricing

More recently, several faster and more accurate valuation methods for CDO valuation have been developed. However, to our knowledge, these approaches cannot readily be applied to CDO^2 valuation.

A fast solution that works for CDO-Squared pricing

In this paper we present a fast accurate numerical method that is applicable to the pricing of both CDO and CDO^2 trades. As well as providing a solution to the CDO^2 problem, it also provides a solution for CDOs that is significantly faster than other methods.

Proxy Integration

Overview

One of the biggest recent developments in credit derivatives has been the growth of the single tranche synthetic CDO market, facilitated by the ever-increasing liquidity of standardised CDX/ITRAXX tranches. Recently, the availability of liquid quotes for these standardised index tranches has enabled market participants to imply a "correlation skew" from the market, creating greater transparency and market consensus on valuation.

The search for yield in a tight-spread environment has also fuelled interest in more complex leveraged structures like CDO of CDOs ("CDO squared" or CDO^2). These new developments have created a need for sophisticated models to generate mark to market and risk in a timely and accurate fashion.

The Gaussian copula has become the market standard model for quoting these structures[1], although there are several other popular variations on this. Traditionally, because of the large numbers of underlying credits involved, these models have been implemented in Monte Carlo. This method has several well-known disadvantages: the risks are inaccurate, and the method requires substantial computational resources.

For transactions with hundreds of credits in the portfolio such as highly diversified CDOs or CDO-squared portfolios, the number of risk factors becomes large. Traditional Monte Carlo implementation becomes impractical. The problem is worsened for supersenior tranches and for highly leveraged trades like CDO-squared where expected losses are low and comparatively few paths in the simulation contribute to the payoff. As a result an extremely large number of paths is required to get reasonable convergence.

More recently, several faster and more accurate valuation methods for CDO valuation have been developed, including Fast Fourier Transform (FFT) and recursion[2]. For the case where all spreads and notionals are equal, reasonable results can also be obtained using the Large Pool Model[3]. However, to our knowledge, these approaches cannot readily be applied to CDO^2 valuation.

In this paper we present a fast accurate numerical method which is applicable to the pricing of both CDO and CDO^2 trades. As well as providing a solution to the CDO^2 problem, it also provides a solution for CDOs that is significantly faster than other methods.

---

[1] David Li, "On Default Correlation: A Copula Approach" Journal of Fixed Income, 9, March 2000, pp 43-54.

[2] Leif Andersen, Jakob Sidenius, Susanta Basu, "All Your Hedges in One Basket" Risk Magazine, November 2003. pp 67-72.

[3] Oldrich Vasicek, "Probability of Loss on Loan Portfolio," KMV Corporation, 1987

Valuation of CDO and CDO^2 in terms of the loss distribution

The valuation of CDO and CDO^2 trades can be conveniently formulated in terms of the portfolio loss distribution.

Consider a CDO tranche on an underlying portfolio of *n* credits, each of which has a notional N_i and recovery R_i. The loss on the portfolio at time *t* is given by:

$$L_t = \sum_i N_i (1 - R_i) 1_{\tau_i \le t}$$

where τ_i denotes the time of default of the i^{th} credit and $1_{\tau_i \le t}$ is the indicator function of default. Then the loss paid out by time *t* on the default leg of the CDO tranche with start loss L_{min} and end loss L_{max} can be written:

$$\hat{L}_t = \min\left(\max\left(L_t - L_{\min}, 0\right), L_{\max} - L_{\min}\right)$$

If we can determine an accurate estimate of the ***loss distribution***, we can work out the value of the trade.

The pricing of a CDO^2 tranche can be formulated in a very similar way, but here if there are *m* underlying CDO tranches, we need to track *m* loss variables $\{L_t^k : k = 1...m\}$, corresponding to each of the underlying portfolios, and work out their ***joint loss distribution*** in order to price the trade.

In the next section we show how to determine an accurate approximation to this joint loss distribution.

An accurate proxy for the joint loss distribution

Uncorrelated Defaults

In the case where all credits are uncorrelated and have equal notional and default probability, we can work out the loss distribution exactly: it is given by the binomial distribution. As an example, consider the case where we have a portfolio of 100 names, all with equal conditional survival probability of 90%.

In Figure 1 we have also plotted a normal distribution with the same mean and variance. It shows visually how well the loss distribution may be approximated by a normal distribution. Although we have made simplifying assumptions of equal notional and survival probabilities here, more generally the overriding principle that ensures that the normal distribution will be a good approximation even when this is no longer true is the central limit theorem, which states that the sum of a set of *finite variance independent* random variables with arbitrary probability distribution converges to the normal distribution as the number of variables increases. In this case the random variables are the indicator functions of default. The accuracy of this approximation is determined by how rapidly the distribution converges to the normal: the graph above indicates that in practice this convergence is fast.

Figure 1. Loss Distribution for Case of Independent Defaults

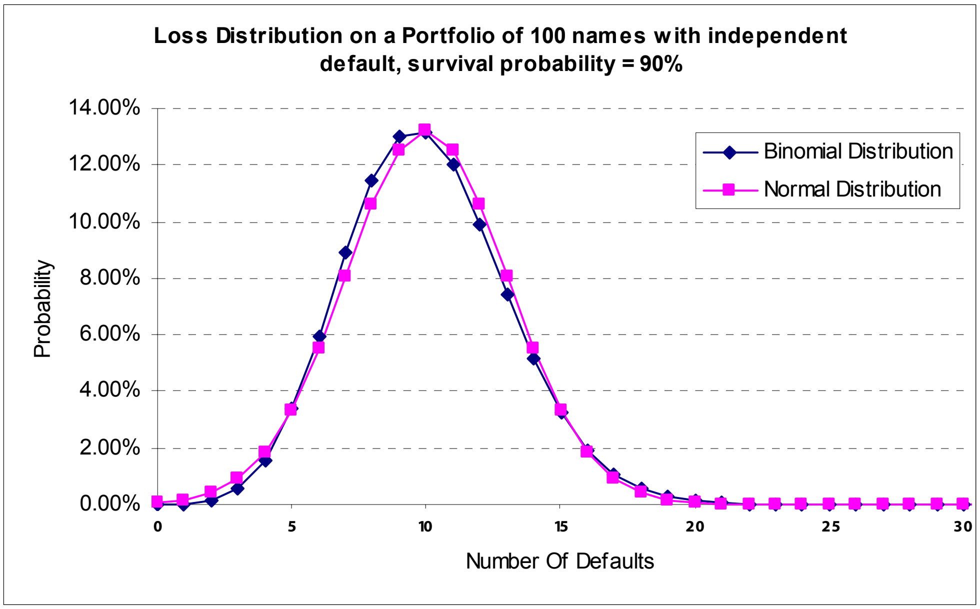

Source: Citigroup.

Correlated defaults

How do we generalise this idea to the case where there is a non-zero correlation between credits? We can use the idea of conditional independence. One of the most popular formulations of the Gaussian copula model is to assume that all credits have the same pairwise correlation. In this case we can write each of the correlated N(0,1) Gaussian variables $\{Z_i\}$ as a linear combination of a central shock variable Y and an idiosyncratic variable $\{X_i\}$: $Z_i = \sqrt{\rho}Y + \sqrt{1-\rho}X_i$ (all variables are independent N(0,1)).

With this choice, conditional on the central shock variable Y all defaults are independent, and hence we can employ an approximation based on the assumption of independence.

Before developing this idea further we note that although we gave this simplified formulation of the Gaussian copula as an example, the approach is not limited to this case. It applies generally to any model in which defaults are independent conditional on a set of factors. It could apply therefore to a Gaussian copula with a more general correlation structure, or a different type of copula such as the t-copula. For this reason we do not give detailed expressions for the conditional survival probabilities: their precise form will depend upon the model used and have been published elsewhere – there are also some powerful techniques for factorisation of general correlation matrices that can be used to cast the problem in this form (Andersen *et al* ibid.).

Assuming that in our chosen model we can calculate the conditional survival probability S_i^Y of credit i at the chosen time-horizon conditioned on a central factor or set of factors Y, we can calculate the distributional parameters of the conditional loss distribution. For generality, we consider the case where there are *m* underlying portfolios defined on the universe of credits and N_i^k represents the notional of credit *i* in portfolio *k*. The loss on portfolio *k* at time *t* is then given by:

$$L_t^k = \sum_i N_i^k (1 - R_i) 1_{\tau_i \le t}$$

We then obtain the following exact expressions for the conditional means and covariances of these loss variables:

$$\mu^k \big| Y = E\big[L_k \big| Y\big] = \sum_i N_i^k (1 - R_i)(1 - S_i^Y)$$

$$Cov\big(L_k, L_{k'} \big| Y\big) = \sum_i N_i^k N_i^{k'} (1 - R_i)^2 \times S_i^Y \big(1 - S_i^Y\big)$$

At this stage we make the key assumption:

Conditional on the central shock factor, the joint distribution of losses is well approximated by a multivariate normal distribution.

Determining the loss distribution using the normal proxy

A CDO or CDO^2 payoff is simply a function of these loss variables at a set of time-horizons. In order to value such a trade based on this approximate form for the loss distribution, we must integrate over the N(0,1) central shock variable Y, and then conditional on each discrete value $Y = Y_i$ approximate the joint loss distribution as described above and form an estimate of the expectation. We provide more detail on exactly how we do this in the next section.

For the example portfolio of 100 names we considered before, we show how the method works when there is a correlation of 25%. Figure 2 shows how integrating over the central shock variable, the loss distributions conditional on different values of Y integrate up to give the overall unconditional loss distribution.

The case of multiple underlying portfolios applicable to CDO^2 trades follows exactly the same pattern. It is worth noting that if there is significant overlap between portfolios, the joint conditional distribution of losses will be highly correlated, even if the defaults are conditionally independent[4].

---

[4] "Managed CDO Squareds", Ratul Roy, Citigroup Global Structured Credit Strategy, June 16, 2004.

Figure 2. Loss Distribution for Case of Correlated Defaults

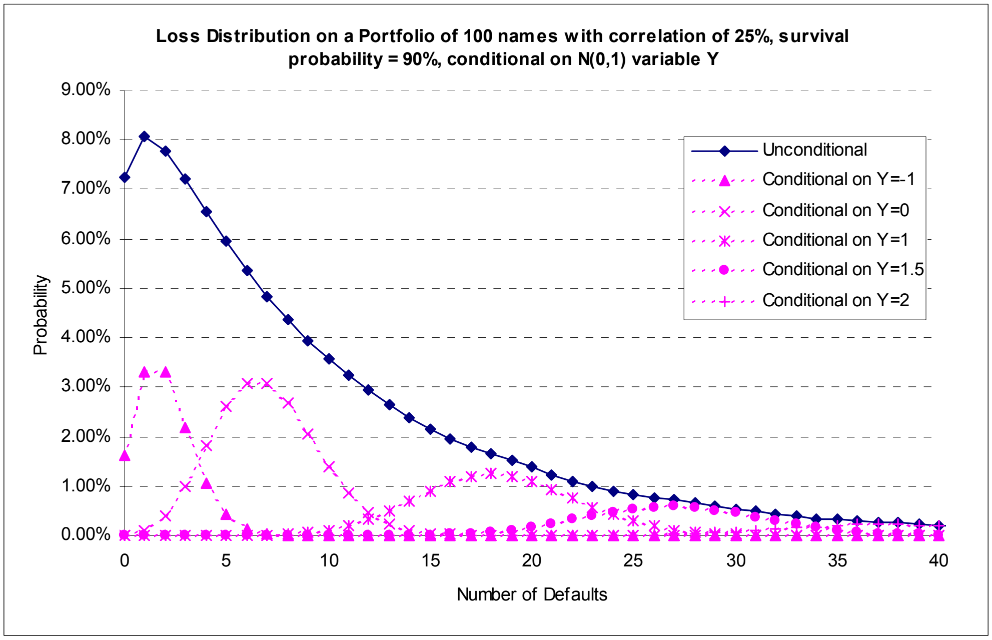

Source: Citigroup.

Determining the expected loss on a tranche using the Proxy Distribution

Having determined the proxy loss distribution we can calculate the conditional expected loss on the tranche.

CDO trades

Recall that for the case of CDOs the loss paid out by time *t* on the default leg of the tranche with start loss L_{min} and end loss L_{max} can be written:

$$\hat{L}_t = \min\left(\max\left(L_t - L_{\min}, 0\right), L_{\max} - L_{\min}\right)$$

This is simply a call spread on the loss variable. Conditional on the central shock variable Y, we assume that the loss L_t is normally distributed. As a result we can derive an analytical expression for the conditional expected loss, which is simply a call-spread on a normally distributed underlying:

$$E\left[\hat{L}_t \middle| Y\right] = (\mu_Y - L_{\min})N\left(\frac{\mu_Y - L_{\min}}{\sigma_Y}\right) + \sigma_Y n\left(\frac{\mu_Y - L_{\min}}{\sigma_Y}\right) - (\mu_Y - L_{\max})N\left(\frac{\mu_Y - L_{\max}}{\sigma_Y}\right) - \sigma_Y n\left(\frac{\mu_Y - L_{\max}}{\sigma_Y}\right)$$

where N denotes the cumulative normal, n the normal density, μ_y the mean and σ_y the standard deviation of the loss conditioned on the central shock variable.

We can therefore derive a very efficient closed-form algorithm for the valuation of CDOs based on discretising the central shock variable Y and numerically integrating these conditional expectations:

$$E[\hat{L}_t] = \int_{-\infty}^{+\infty} n(Y) E\left[\hat{L}_t \middle| Y\right] dY \approx \sum_{\{Y_j\}} w_j E\left[\hat{L}_t \middle| Y_j\right]$$

where $\{Y_j\}$ is the chosen discretisation of the Y axis and w_j a weighting factor dependent upon the chosen numerical integration method. Good results can be obtained with a very simple algorithm like the trapezium rule.

CDO^2 trades

We also comment briefly on the case of CDO^2 trades. In this case the dimensionality of the problem is of size *m*: number of underlying portfolios. Usually this number is larger than 3, so there is no simple closed form expression for the conditional expected loss. The problem then becomes one of estimating a multivariate normal integral. One way of efficiently achieving this is by using an *m*-dimensional Monte Carlo integration. We note that this Monte Carlo is vastly lower in variance than the standard implementation of the Gaussian copula because the number of factors in the Monte Carlo is much lower and the variables that are simulated are smooth loss variables, not Bernouilli indicators of default. In the standard approach, the Bernouilli indicators and hence the portfolio loss are inaccurately reproduced because only a small percentage of paths result in default events.

Results

We find that proxy integration generates accurate PV and risks. The performance overheads are very low compared to standard Monte Carlo.

Spread Risk

In Figure 3 we show some example Spread01 numbers for a CDO trade compared against first principles Monte Carlo, based on a spread bump of 10bp. We plot the by name Spread01 against the spread for each underlying name which shows very visually the accuracy of the method.

Figure 3. Spread Risk for CDO Trade

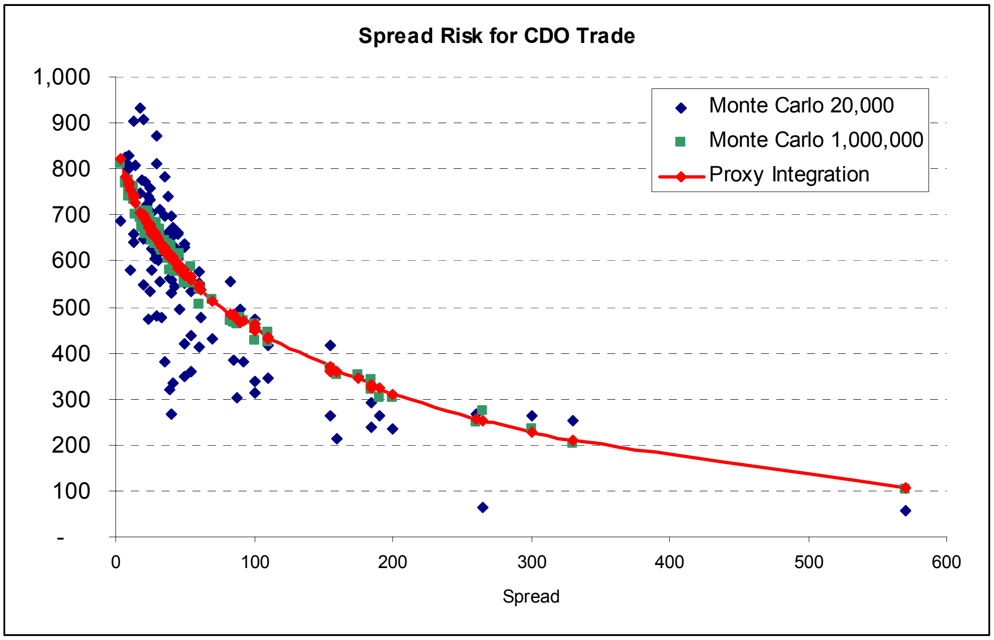

Source: Citigroup.

Figure 4. Spread Risk for CDO^2 Trade

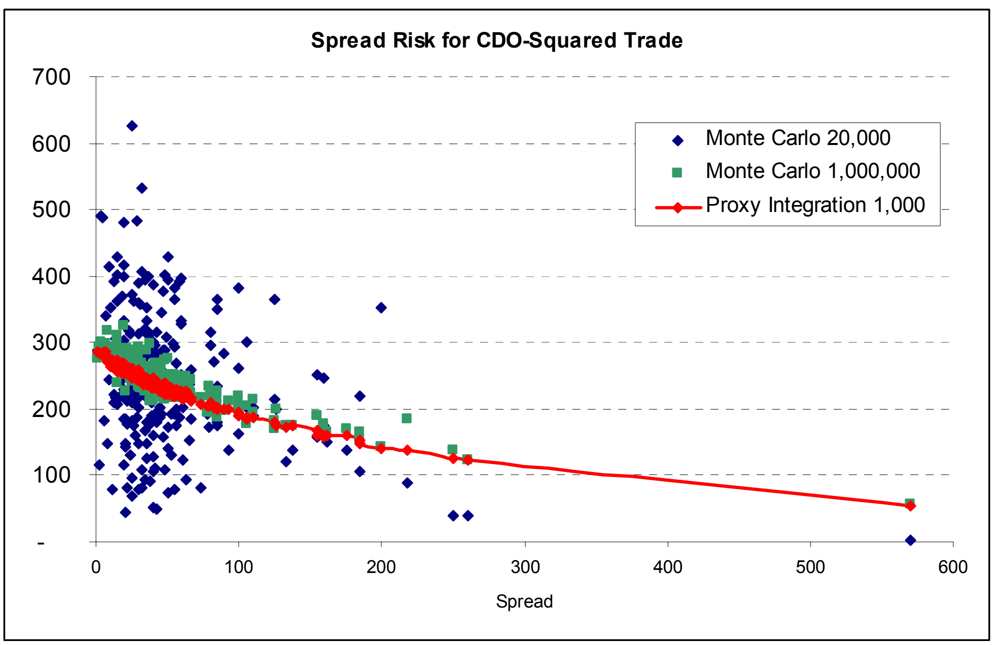

Source: Citigroup.

The first example is a mezzanine CDO tranche with a maturity of 4.5 years and 100 underlying credits. We assume a correlation of 25%.

The second example is a CDO^2 tranche consisting of 6 underlying CDO mezzanine tranches and 225 underlying credits and has a 4 year maturity. For the case of CDO^2 trades, the valuation in this implementation still involves a Monte Carlo procedure.

However, as can be seen from Figure 4, with the standard Monte Carlo we must perform a 225-dimensional simulation with 1 million paths, whereas in Proxy Integration we only need to do a 6 dimensional simulation with 1000 paths to get somewhat better accuracy. It is the fact that proxy integration requires far fewer paths to get a good level of accuracy and is of much lower dimension that makes it such a powerful technique.

Accuracy

In general, we find empirically by testing across many CDO and CDO^2 trades that the proxy integration method is capable of giving results for the breakeven spread of CDO and CDO^2 with a relative difference of at most a few percent from the values obtained with standard Monte Carlo with a million paths, which amounts in absolute terms to at most a few basis points and usually sub-basis point accuracy

As described above the risks are also much more accurately reproduced. As a result it is possible to use much smaller spread bumps of less than 1bp and still get excellent results for the Spread01.

Performance

We give performance results for these 2 examples priced on a single processor 2.66GHz Pentium 4 desktop. Because of efficiencies in the algorithm the time taken for the risk numbers is not a high multiple of the time for a single valuation. Clearly the calculation times are far lower than those required in a standard Monte Carlo to get accurate results:

	Proxy Integration (1,000 paths for CDO^2, closed form for CDO)		Monte Carlo (1,000,000 paths)	
	PV	PV and Spread01	PV	PV and Spread01
CDO	0.04 sec	0.16 sec	220 sec	350 sec
CDO^2	0.4 sec	1.0 sec	800 sec	1150 sec

The Spread01 timings correspond to calculating spread risk for every underlying name in each trade: 120 risks in the case of the CDO and 225 risks in the case of the CDO^2.

Conclusions

The proxy integration method described in this research paper provides a very efficient and accurate way of calculating prices and risks of CDO and CDO^2 trades. The major advantages of this method are:

- CDO pricing is faster than other closed form techniques.
- CDO^2 pricing is much faster and more accurate than the only other published technique (Monte Carlo).
- The method does not have any problems with diverse spreads, notionals and recoveries – these can cause difficulties for methods such as FFT.

Although the method involves one key approximation: that the conditional joint loss distribution is well approximated by the multivariate normal, we find empirically that with real trades and market data we get very satisfactory accuracy.

There are also several possible generalisations of the work described in this paper, most notably some significant improvements in the speed of calculation of the Greeks.

We believe that the work presented here will contribute towards helping market participants to get more transparent and accurate representations of the risk of their structured credit positions. From a world of arduous Monte Carlo simulation, we hope that this work will encourage market participants to get “back to normal”.

Notes

Notes

Notes

Disclosure Appendix

ANALYST CERTIFICATION

I, David Shelton, hereby certify that all of the views expressed in this research report accurately reflect my personal views about any and all of the subject issuer(s) or securities. I also certify that no part of my compensation was, is, or will be directly or indirectly related to the specific recommendation(s) or views in this report.

Other Disclosures

ADDITIONAL INFORMATION AVAILABLE UPON REQUEST

Citibank, N.A., London Branch and Citigroup Global Markets Inc, including its parent, subsidiaries, divisions and/or affiliates ("the Firm"), may make a market in the securities discussed in this report and may sell to or buy from customers, as principal, securities recommended in this report. The Firm may have a position in securities or options of any issuer recommended in this report. The Firm may be regular issuers of, and trade in (including position taking), financial instruments linked to securities, which may have been reported on in this research report. The Firm may perform or solicit investment banking or other services from any issuer recommended in this report. An employee of the Firm may be a director of an issuer recommended in this report. Within the past three years, the Firm may have acted as manager or co-manager of a public offering of the securities of any issuer recommended in this report. Securities recommended, offered, or sold by the Firm : (i) are not insured by the Federal Deposit Insurance Corporation; (ii) are not deposits or other obligations of any insured depository institution (including Citibank); and (iii) are subject to investment risks, including the possible loss of the principal amount invested. Past performance is not a guarantee of future results. This report does not take into account the investment objectives, financial situation or particular needs of any particular person. Investors should obtain advice based on their own individual circumstances before making an investment decision. Investing in non-U.S. securities, including ADR's entails certain risks. The securities of non-U.S. issuers may not be registered with, nor be subject to the reporting requirements of, the U.S. Securities and Exchange Commission. There may be limited information available on foreign securities. Foreign companies are generally not subject to uniform audit and reporting standards, practices and requirements comparable to those in the U.S. Securities of some foreign companies may be less liquid and their prices more volatile than securities of comparable U.S. companies. In addition, exchange rate movements may have an adverse effect on the value of an investment in a foreign securities and its corresponding dividend payment for U.S. investors. Net dividends to ADR investors are estimated, using withholding tax rates conventions, deemed accurate, but investors are urged to consult their tax advisor for exact dividend computations. Although information has been obtained from and is based upon sources the Firm believes to be reliable, we do not guarantee its accuracy and it may be incomplete or condensed. All opinions and estimates constitute the Firm 's judgement as of the date of the report and are subject to change without notice. This report is for informational purposes only and is not intended as an offer or solicitation for the purchase or sale of a security. This research report does not constitute an offer of securities. Any decision to purchase securities mentioned in this research must take into account existing public information on such security or any registered prospectus. Investing in non-US securities by US persons may entail certain risks. Investors who have received this report from the Firm may be prohibited in certain US States from purchasing securities mentioned in this report from the Firm; please ask your Financial Consultant for additional details. This report is distributed in the United Kingdom by Citibank, N.A. London Branch or Citigroup Global Markets Limited, Citigroup Centre, Canada Square, Canary Wharf, London E14 5LB, UK. This material is directed exclusively at market professional and institutional investor customers in the United Kingdom and is not for distribution to private customers in the United Kingdom, as defined by the rules of the Financial Services Authority, who should not rely on this material. Moreover, any investment or service to which the material may relate will not be made available to such private customers. This material may relate to investments or services of a person outside of the United Kingdom or to other matters which are not regulated by the Financial Services Authority and further details as to where this may be the case are available upon request in respect of this material. If this publication is being made available in certain provinces of Canada by Citigroup Global Markets (Canada) Inc. ("The Firm Canada"), the Firm Canada has approved this publication. If this report was prepared by the Firm (excluding Nikko Citigroup Limited) and distributed in Japan by Nikko Citigroup Limited, it is being so distributed under license. This report is made available in Australia, to non-retail clients through Citigroup Global Markets Australia Pty Limited (ABN 64 003 114 832 and AFSL No. 240992), a participant of the ASX Group and a participant of the Sydney Futures Exchange Limited and to retail clients through Smith Barney Citigroup Australia Pty Ltd (ABN 10 009 145 555 and AFSL No. 240813), a participant of the ASX Group. In New Zealand it is made available through Citigroup Global Markets New Zealand Limited, a member firm of the New Zealand Stock Exchange. Citigroup Global Markets (Pty) Limited is incorporated in the Republic of South Africa (company registration number 2000/025866/07) and its registered office is at Citibank Plaza, 145 West Street, Sandown, Sandton, 2196, Republic of South Africa. The investments and services contained herein are not available to private customers in South Africa. This publication is made available in Singapore through Citigroup Global Markets Singapore Pte Ltd, a Capital Markets Services license holder. This report is being distributed in Hong Kong by or on behalf of, and is attributable to Citigroup Global Markets Asia Limited, 20th Floor, Three Exchange Square, Hong Kong. Citigroup Global Markets Inc. is a member of the Securities Investor Protection Corporation (SIPC).