SOUTH AFRICAN ACTUARIAL JOURNAL
SAAJ **14** (2014) 143–57

ABSTRACTS OF RECENT POSTGRADUATE THESES AND DISSERTATIONS AT SOUTH AFRICAN UNIVERSITIES

The impact of firm-specific factors on the cross-sectional variation in Johannesburg-Securities-Exchange-listed equity returns

By JD van Heerden for PhD at the University of Stellenbosch, 2014

The aim of this study is to examine the impact of technical and fundamental—referred to as firm-specific—factors on the cross-sectional variation in equity returns on the Johannesburg Securities Exchange (JSE). A secondary objective is to examine the effect that time, holding, or payoff, period and liquidity may have on the identity and explanatory power of factors on the cross-section of equity returns. A cross-sectional regression approach, a factor-portfolio approach and an extreme-performer approach are followed to reach the objectives and to allow for comparison and robustness-tests to be performed on the JSE. Monthly data were collected for all identified factors through numerous data sources on all shares listed on the JSE for the period January 1994 to May 2011. The data have been prepared to correct for potential statistical biases that may affect the results, including data snooping, infrequent trading, survivorship bias, look-ahead bias and outliers. This lengthy period allows for the formation of two independent subsamples, each covering a full investment cycle, enabling in-sample and out-of-sample empirical research and testing to be conducted. The results of the three approaches are strongly correlated. It suggests that a strong value effect is present and robust, while a size effect appears to be sensitive to time, liquidity and payoff period. A price reversal effect is observed over both a short term—represented by prior one-month returns—and longer term—prior 60-month returns—but this observation is dependent on the subsample, level of sample liquidity and payoff period. A momentum effect, based on prior 6- and 12-month returns, appears to be robust. Multifactor analyses show that value and momentum factors are collectively significant in explaining the cross-section of returns across all time periods and level of liquidity.

The excess returns obtained through constructing portfolios based on the identified

factors cannot be explained by known market models. Hence, the findings of this thesis lead to the rejection of the efficient-market-and-CAPM joint hypothesis within the South African context, and suggest that market anomalies exist or that current market models are incorrectly specified for the JSE.

Pairs trading: a copula approach

By CM Augustine for MPhil at the University of Cape Town, 2014

Pairs trading is an arbitrage strategy that involves identifying two pairs of stocks known to move together historically and trading on them when relative mispricing occurs. The strategy involves shorting the overvalued stock and simultaneously going long on the undervalued stock and closing the positions once the prices have returned to fair values. The cointegration method and the distance method are the most common techniques used in pairs-trading strategy. However, these strategies rely on normality of the distribution of what is often known as the spread, which is the measure of divergence between the two stocks. This assumption is one of the main drawbacks of these strategies. It may lead to missed or inaccurate trading signals. The purpose of this dissertation is to explore an alternative approach to pairs trading by use of copulas. This dissertation aims to investigate if copulas can improve the profitability of pairs trading. To achieve this aim, results of pairs trading by use of copulas are compared against those of cointegration and distance methods.

The recovery theorem: expounded and applied

By A Backwell for MPhil at the University of Cape Town, 2014

This dissertation is concerned with Ross's recovery theorem. It is generally held that a forward-looking probability distribution is unobtainable from derivative prices, because the market's risk-preferences are conceptually inextricable from the implied real-world distribution. Ross's result recovers this distribution without making the strong preference assumptions assumed necessary under the conventional paradigm. This dissertation aims to give the reader a thorough understanding of Ross Recovery, both from a theoretical and practical point of view. It starts with a formal delineation of the model and proof of the central result, motivated by the informal nature of Ross's working paper. This dissertation relaxes one of Ross' assumptions and arrives at the equivalent conclusion. This is followed by a critique of the model and assumptions. An a priori discussion only goes so far, but potentially problematic assumptions are identified, chief amongst which being time-additive preferences of a representative agent. Attention is then turned to practical application of the theorem. The author identifies a number of obstacles to applying the result—some of which are somewhat atypical and have not been directly addressed in the literature—and suggests potential solutions. A salient obstacle is calibrating a state price matrix. This leads to an implementation of Ross recovery on the FTSE/JSE Top40. The suggested approach is found to be workable, though certainly not the final word on the matter. A testing framework for the model is

discussed and the dissertation is concluded with a consideration of the findings and the theorem's applicability.

Robust portfolio construction: controlling the alpha-weight angle
by G Bailey for MPhil at the University of Cape Town, 2013

Estimation risk is widely seen to have a significant impact on mean-variance portfolios and is one of the major reasons the standard Markowitz theory has been criticised in practice. While several attempts to incorporate estimation risk have been considered in the past, the approach of Golts & Jones represents an innovative approach to incorporate estimation risk in the sample estimates of the input returns and covariance matrix. In this project, the author discusses the theory introduced by Golts & Jones (2009) which looks at the direction and the magnitude of the vector of optimal weight and investigates them separately, with focus on the former. The author demystifies the theory of Golts & Jones with focus on both mathematical reasoning and practical application. The author shows that the distortions of the mean-variance optimisation process can be quantified by considering the angle between the vector of expected returns and the vector of optimised portfolio positions. Golts & Jones call this the alpha-weight angle. The author shows how to control this angle by employing robust optimisation techniques, which are also explored as a main focus in this project. The author applies this theory to the South African market and shows that one can indeed obtain portfolios with lower risk statistics, especially in times of economic crisis.

Constructing volatility surfaces for managed funds
By T Brinkmann for MPhil at the University of Cape Town, 2014

In this dissertation, a methodology is developed for constructing a volatility surface for a managed fund by extending the work of Bakshi et al. and Taylor. The power utility assumption—with constant relative risk aversion for a specific maturity—and historical returns series data are used for the identified factors influencing the return of the fund and the fund itself. The coefficient of relative risk aversion for a specific maturity and market is estimated from quoted option prices on a market index. This is used in combination with the identified factors and fund return series to estimate the risk-neutral skewness of the fund. An optimisation procedure is then used to determine the volatility smile of the fund for a specific maturity. Thereafter, the volatility surface of the fund is constructed by repeating each step for different maturities. Although this methodology produces sensible results, the optimisation routine used is sensitive to initial values and constraints.

A utility-driven change of measure
By K Carolissen for MPhil at the University of Cape Town, 2014

The author demonstrates how a change of probability measure can be carried out based on the risk preference of a representative investor. Using the stochastic discount

factor and the Radon-Nikodym derivative, the author is able to obtain the risk-neutral measure given a real-world measure and a preference structure defined by a utility function. This methodology is then used to attribute the sources of skewness in the risk-neutral measure.

Modelling illiquid volatility skews

By SM Crowther for MPhil at the University of Cape Town, 2014

Most markets trade liquidly in options on the market index, in fact they often trade at a wide range of strike levels. Thus, using the Black–Scholes model, one can obtain the implied volatilities at the various strike levels, forming the associated implied volatility skew of the respective market under consideration. This, however, is not always feasible when it comes to the individual stocks within the market, as single stock options trade much less frequently. This dissertation makes use of data from the Eurozone, in particular the author considers the Euro Stoxx 50 market index and its underlying constituents. Options written on the Euro Stoxx 50 and its constituents are highly liquid, and volatility skews are obtained for the market as well as for most of the single stocks within the market. The author then artificially created three cases of illiquid markets, each with increasing degrees of sparseness mimicking various possible realities. Using principal-component analysis, this dissertation aims to find an appropriate model for relating the volatility skew of the index to that of single stocks within the market in order to fill gaps in the data of the skews of the individual stocks. Results indicate that simpler models perform similarly in all scenarios of sparseness whereas the performance of more complex models decrease as the data becomes sparser. This indicates that basic relationships can be formed between the index and single stocks in cases with relatively low levels of trade in the market but more accurate estimates are more difficult to achieve. However, if one uses the skew data, as is, as an input to the models, their performance remains broadly the same using the full data set and using monthly information. This is encouraging, as it means one can fill gaps in the individual stocks' skew data with as good a fit as if one modelled with a full set of data.

A post-crisis investigation into the performance of GARCH-based historical and analytical value at risk on the FTSE/JSE

by A de Alessi for MPhil at the University of Cape Town, 2013

This paper is an investigation into the performance of GARCH-based value-at-risk (VaR) models on the South African FTSE/JSE Top 40 Index. Specifically, this paper investigates whether stability has returned to the VaR measure following its poor performance during the latest global financial crisis (2007). GARCH models are used in both an analytic and historical approach for modelling 1%, 2,5% and 5% daily VaR for a three-year backtest period (2010–2012). Four distributions are used: the normal, generalised error, *t* distribution and the skewed *t* distribution. A particular question asked by this paper, is whether the data from the latest financial crisis (2007) should be used

in estimating VaR in a post-crisis market. To investigate this, all models are re-estimated using data that has the financial crisis or high volatility period removed, and the results across the two data sets are compared. The take-away point from this research is that the volatility-clustering mechanism inherent in every GARCH model is capable of producing accurate VaR estimates in a post-downturn or lower-volatility market even when the data on which the model was estimated contain financial downturn or volatile data. There is strong evidence suggesting stability has returned to this measure however caution remains over using over-simplified models.

Inflation modelling for long-term liability-driven investments

By J de Kock for MPhil at the University of Cape Town, 2014

A regime-switching model allows a process to switch randomly between different regimes which have different parameter estimates. This study investigates the use of a two-regime-switching model for inflation in South Africa as a means of hedging inflation-linked liabilities of a financial institution. Each regime is modelled using an autoregressive process with different parameters and the change in regimes is governed by a two-state Markov chain. Once the parameters have been estimated, the predictive validity of the regime-switching process as a model for inflation in South Africa is tested and a hedging strategy is outlined for a set of inflation-linked cashflows. This hedging strategy is determined through the use of rand-per-point methodology that is applied to the inflation-linked cashflows and inflation-linked bonds. This hedging strategy was shown to be profitable over the period from January 2008 to June 2013.

Statistical arbitrage in South Africa

By J-J D de Wit for MPhil at the University of Cape Town, 2014

This study investigates the performance of a statistical arbitrage portfolio in the South African equity markets. A portfolio of liquid stock pairs that exhibit cointegration is traded for a ten-year period between the years 2003 and 2013. Without transaction costs, the portfolio has an encouraging Sharpe ratio of 2,1. When realistic transaction costs are factored in, the Sharpe ratio drops to 0,43. The results underline the theoretical profitability of statistical arbitrage as a trading strategy and highlight the importance of transaction costs in a real-world setting.

The benefits of a tree-based model for stock selection in a South African context

By M Giuricich for MPhil at the University of Cape Town, 2014

Quantitative investment practitioners typically model the performance of a stock relative to its benchmark and the stock's fundamental factors in a classical linear framework. However, these models have been found empirically to be unsuitable for capturing higher-order relationships between a stock's return relative to a benchmark and its fundamental factors. This dissertation studies the use of classification and regression tree (CART) models for stock selection within the South African context, with the focus

being on the period from when the global financial crisis began in early 2007 until December 2012. By utilising four types of portfolios, a CART model is directly compared against two traditional linear models. It is seen that during the period focused on, the portfolios based on the CART model deliver the best excess return and risk-adjusted return, albeit in most cases modestly above the returns delivered by the portfolios based on the linear models. This is observed in the hedge-fund style and long-only portfolios constructed. Moreover, it is observed that the CART-based portfolios' returns are not correlated with those from the linear-model-based portfolios. This observation suggests that CART models offer an attractive option to diversify model risk within the South African context.

Hedge fund of funds investment process

By NM Hossain for MPhil at the University of Cape Town, 2014

The objective of this dissertation is to develop and test an investment process for hedge fund of funds (HFoFs) in South Africa. The dissertation proposes a three-tiered process, adapted from the works of Lo.

Step one of the process involves the categorisation of hedge funds into broadly defined groups based on predefined factors. Two classification methodologies are examined herein to determine optimal category definitions. These are 1) an adaption of the classification developed by Schneeweis & Spurgin, based on the correlation of hedge funds to an appropriate benchmark and the returns offered by these hedge funds, and 2) classification by cluster analysis. Once a finite set of classification is defined, step two of the process uses a minimum-variance optimisation, based on forward-looking parameter estimates of return and covariance to compute the optimal capital allocation to these categories. The final stage of the process employs a mixture of quantitative and qualitative analysis to allocate capital within categories to individual hedge funds.

Robust portfolio construction using sorting signatures

By L Kasenene for MPhil at the University of Cape Town, 2014

Mean-variance analysis—introduced by Harry Markowitz—has been criticised in the past mainly due to the counterintuitive and unstable nature of the resultant portfolios from the optimisation. These disappointing results have been linked to the presence of estimation error in the estimates of the expected returns and covariances, which serve as inputs to the optimisation. Several attempts have been made to produce more reliable estimates, with a significant amount of effort and resources placed on the estimation of expected returns, which is generally a more difficult task than the estimation of covariances. Almgren and Chriss provide a methodology for portfolio selection in which the order of expected returns replaces the numerical values of the returns. This framework allows full use of the covariance matrix, in a method analogous to mean-variance optimisation. The author adopts this framework in the analysis together with the robust optimisation technique introduced by Golts and Jones, which improves the

estimate of the covariance matrix by direct modification in the optimisation process. Golts and Jones argue that a reduction of the angle between the input return forecasts and the output portfolio positions results in more investment relevant portfolios, in line with the investment manager's insights. They relate this angle to the condition number of the covariance matrix and use robust optimisation to improve the conditioning of this matrix. Assuming perfect alpha foresight of an investment manager, the author applies a combination of the techniques of Almgren and Chriss and Golts and Jones to South African equity data and shows that the resultant robust portfolios, though conservative in their risk-adjusted return statistics, are more diversified and exhibit lower leverage than mean-variance portfolios. The author further shows that independent of the optimisation method, there is a marginal difference in the performance of portfolios created using ordering information and actual returns.

Realised volatility estimators

By S Königkrämer for MPhil at the University of Cape Town, 2014

This dissertation is an investigation into realised volatility (RV) estimators. Here, RV is defined as the sum of squared returns (SSR) and is a proxy for integrated volatility (IV), which is unobservable. The study focuses on a subset of the universe of RV estimators. The author examines three categories of estimators: historical, high frequency (HF) and implied. The need to estimate RV is predominantly in the hedging of options and is not concerned with speculation or forecasting. The main research questions are: (1) What is the best RV estimator in a historical study of S&P 500 data? (2) What is the best RV estimator in a Monte Carlo simulation when delta hedging synthetic options? (3) Do the findings support the stylised fact of asymmetry in time scales? In the answering of these questions, further avenues of investigation are explored. Firstly, the VIX is used as the implied volatility. Secondly, the Monte Carlo simulation generates stock price paths with random components in the stock price and the volatility at each time point. The distribution of the input volatility is varied. The question of asymmetry in time scales is addressed by varying the term and frequency of historical data. The results of the historical and Monte Carlo simulation are compared. The SSR and two of the HF estimators perform best in both cases. Accuracy of estimators using long term data is shown to perform very poorly.

Estimating credit default swap spreads from equity data

By J Kooverjee for MPhil at the University of Cape Town, 2014

Corporate bonds are an attractive form of investment as they provide higher returns than government bonds. This increase in returns is usually associated with an increase in risk. These risks include liquidity, market and credit risk. This dissertation focuses on the modelling of a corporate bond's credit risk by considering how to estimate the credit default swap spread of a firm's bond. A structural credit model is used to do this. In this dissertation, the author implements an extension of Merton's model which is based on

the use of the implied volatilities of options on the company's stock to estimate model parameters. Such an approach is interesting as it provides an insight into the linkages between credit markets and options markets.

Volatility derivatives in the Heston framework

By H Kriel for MPhil at the University of Cape Town, 2014

A volatility derivative is a financial contract where the payoff depends on the realised variance of a specified asset's returns. As volatility is in reality a stochastic variable, not deterministic as assumed in the Black–Scholes model, market participants may surely find volatility derivatives to be useful for hedging and speculation purposes. This study explores the construction and calibration of the Heston stochastic-volatility model and the pricing of some volatility derivatives within this framework.

Hawkes processes and some financial applications

by B Lapham for MBusSc at the University of Cape Town, 2014

The self-exciting point process, which is now more commonly known as the Hawkes process, is a model for a point process on the real line introduced by Hawkes. The distinguishing feature of such processes is that they allow all past events to affect the intensity function at the current time.

Over the years such processes have been applied in seismology and neurophysiology in particular, and in more recent years there have been significant financial applications. In almost all of these applications, the route used to find the maximum-likelihood estimates (MLEs) is direct numerical maximisation (DNM) of the likelihood. An EM algorithm, which makes use of the Poisson cluster process interpretation of the Hawkes process, is an alternative route to the MLEs. This particular EM algorithm has received attention in the literature and has been claimed to have advantages over DNM of the likelihood. The author carries out a simulation study for a simple Hawkes process to clarify statements made in the literature about these advantages. For the simple Hawkes process models that the author considers, DNM of the likelihood is the preferable route to finding the MLEs.

The author then uses DNM of the likelihood to fit marked Hawkes-process models to South African asset data. These applications to South African data include the modelling of extreme asset returns and the forecasting of conditional value at risk (VaR) and expected shortfall (ES). The models investigated include mostly models found in the literature, but also include some variations introduced in the dissertation. In a backtesting exercise, the author compares the conditional VaR and ES forecasts found by using the marked Hawkes-process models with those found via some non-standard stochastic-volatility (SV) models. The author finds that the marked Hawkes-process models give mostly competitive forecasts of conditional VaR and ES when compared with the non-standard SV models.

Modelling covariates of infant and child mortality in Malawi

By C Lemani for MPhil at the University of Cape Town, 2014

Mortality of children under the age of five has been the main target of public health policies. There has been a significant decline in under-five mortality in the twentieth century in almost all countries regardless of initial levels and socio-economic factors, although the rate of decline has been different in different regions. Malawi, a country in the sub-Saharan region, is characterised by high infant and child mortality. Using data from the 2010 Malawi Demographic and Health Survey, infant mortality in Malawi was estimated at 66 deaths per 1000 births while child mortality was at 50 deaths per 1000 births.

Studies have been conducted to identify covariates of infant and child mortality in Malawi but none of these used recent data and none has included HIV/AIDS as a risk factor. This study aims at examining bio-demographic, socio-economic and environmental factors associated with infant and child mortality in Malawi. Malawi Demographic and Health Survey (DHS) data for 2004 and 2010 are used. Two methods of analysis are applied to both 2004 DHS and 2010 DHS. These are the logistic regression method and survival-analysis method.

The results of the study show that HIV status of the mother and length of the preceding-birth interval were significantly associated with both infant and child mortality during both time periods. Other significant covariates include birth order, age of the mother at birth of the child, sex of the child, education of the mother and father, and wealth index. The risk of death among infants born to HIV-positive mothers declined during the inter-survey period but increased for children aged 1 to 5. Scaling of Prevention of Mother to Child Transmission (PMTCT) programmes, raising awareness on the importance of long preceding-birth interval and advocating parents' education will therefore help in reducing childhood deaths among others.

An evaluation of the proposed IFRS 4 Phase II measurement methodology: the impact on South African life insurers

By S Marszalek for MBusSc at the University of Cape Town, 2014

Nearly 20 years after inception, the insurance accounting project of the International Accounting Standard Board (IASB) is nearing completion. The June 2013 International Financial Reporting Standard 4 (IFRS 4) Exposure Draft represents a likely picture of the future of global insurance financial reporting and it is important that insurers begin to understand and prepare for the changes it will bring. This dissertation explores the key principles and likely impacts of the IFRS 4 Phase II standard, in its current proposed form, in the South African life insurance context. In particular, the proposed IFRS 4 Phase II approach to profit reporting is contrasted with the current financial-soundness valuation (FSV) approach for simple illustrative term- and endowment-assurance products. The results of this comparison form the basis for a discussion of the impacts which the new profit-reporting standard will have on insurance-contract liabilities and hence profit profiles over time, and an assessment of whether the changes embodied in the new standard better meet the objectives of insurance financial reporting and the needs

of the users of financial statements. This dissertation focuses on key areas where there is a high degree of certainty in the exposure draft, and touches more lightly on those areas where change is still expected. The findings indicate that IFRS 4 will result in insurer financial reporting being more principles-based, better meeting the requirements of fundamental financial-reporting principles and being more comparable with insurer financial reporting internationally. These findings support the conclusion that a move to IFRS 4 for insurer financial reporting in South African will be beneficial to users of financial statements in making economic decisions.

Efficient implementation of Heston-Hull & White model

By S Maze for MPhil at the University of Cape Town, 2014

A model with a stochastic interest-rate process correlated to a stochastic volatility process is needed to accurately price long-dated contingent claims. Such a model should also price claims efficiently in order to allow for fast calibration. This dissertation explores the approximations for the characteristic function of the Heston-Hull & White model introduced by Grzelak and Oosterlee. Fourier cosine expansion pricing is then used to price contingent claims under this model, which is implemented in MATLAB. The author finds that the model is efficient, accurate and has a relatively simple calibration procedure. In back-tests, it is determined that the Heston-Hull & White model produces better hedging profit and loss results than a Heston or a Black and Scholes model.

A multi-regional adaptation of the ASSA2008 AIDS and Demographic Model for population projections

By GT Mofokeng for MPhil at the University of Cape Town, 2014

The traditional models used in population projections rely on the net-migration method. The ASSA2008 AIDS and Demographic model is one such model.

In this research, the nine provinces and the four population groups are aggregated to give rise to three regions. Using STATA12, the directional migration tables for the years covering the period 1996–2007 between the three regions, by age and sex, based on a 10% sample of the 2001 Census and a 2,5% sample of the 2007 Community Survey, are produced.

Using MATLAB 2011a with built-in Levenberg–Marquardt algorithm with nonlinear least-squares methods, Rogers–Castro multi-exponential age schedules are fitted to the census or survey migration data in order to obtain parameters used to estimate migration rates in the model for the period 1996–2007. After 2007, migration rates are extrapolated roughly linearly, assuming that migration will trend towards zero over a fixed number of years.

The multi-regional adaptation of the ASSA model is tested and found to work, with a minor re-calibration to the HIV data for 2008.

The projected regional-population age structure and size implied by the model for 1996–2025 are consistent with the same estimates implied by the net migration model,

and so are the projected net-migration rates per 1 000. The level of the migration rates assumed in the multi-regional model accounts for an average of 89% of the change in the estimates of the population size relative to those generated by the net-migration model, and the use of multi-regional modelling itself accounts for 11% of these changes.

The proportions of the changes attributable to the level of migration rates assumed in the multi-regional model, and the use of the multi-regional modelling, show that the choice of the method by which population projections are done is important.

Finally, the three-region model can be extended to a nine-province model that recognises that each province has unique demographic dynamics, but the construction of such a model requires a significant amount of extra work due to its size and complexity.

The two-dimensional COS method for pricing early-exercise and discrete barrier options under the Heston model

By R Moir for MPhil at the University of Cape Town, 2014

The author focuses on the pricing of Bermudan and barrier options under the dynamics of the Heston stochastic volatility model. The two-dimensional nature of the Heston model makes the pricing of these options problematic, as the risk-neutral expectations need to be calculated at each exercise and observation date along a continuum of the two state spaces. We examine the 2D-COS method, which makes use of Fourier cosine expansions in each of the two dimensions in order to approximate the integrals. Using the fast Fourier transform, the author is able to calculate the cosine series coefficients efficiently at each exercise and observation date. A construction of this method is provided and the author conducts numerical experiments to evaluate its speed and accuracy.

Investigating biases in census questions on mortality using Agincourt Health and Demographic Surveillance System data

By PM Nduru for MPhil at the University of Cape Town, 2014

An understanding of the errors found in census questions is important in order to assess the level of confidence in the census data and to get an appreciation of the impact of using these data in estimating mortality derived from census data. While demographic methods are often used to determine the direction of bias in the data, direct evaluation studies are required to determine the nature and extent of biases with more accuracy. Equally important is an understanding of the characteristics of the respondents who produce better responses in censuses. This can be used to inform selection of better respondents in order to improve the quality of the collected data.

This research uses census data from a survey, which is matched to the longitudinal Health Demographic Surveillance System site (HDSS) data from Agincourt, Limpopo Province, South Africa, in order to assess the biases found in data used in child and adult mortality estimation that uses indirect techniques. The research also assesses the accuracy of data used to measure mortality directly, and assesses the potential of additional questions aimed at cause of death analysis and completeness of death

registration in censuses. This is done by comparing the census information to the true record as captured by the HDSS.

The results from this research show that indirect child mortality estimates are less robust than indirect adult mortality estimates because of the relatively poor quality of responses to the children ever born or children surviving questions compared to the orphanhood questions. Questions on household deaths in the year before the survey and questions on registered deaths have the potential to allow for the estimation of recent mortality levels at all ages and the completeness of death registration. However, there is a need to train the census enumerators to be able to elicit accurate information, in particular on ages at death. In addition, young (aged 15–29) adult female respondents generally give more accurate responses to all the questions in the census.

Optimal tree methods

By R Rudd for MPhil at the University of Cape Town, 2014

Although traditional tree methods are the simplest numerical methods for option pricing, much work remains to be done regarding their optimal parameterisation and construction. This work examines the parameterisation of traditional tree methods as well as the techniques commonly used to accelerate their convergence. The performance of selected, accelerated binomial and trinomial trees is then compared to an advanced tree method, Figlewski and Gao's adaptive mesh model, when pricing an American put and a down-and-out barrier option.

An assessment of the application of cluster-analysis techniques to the Johannesburg Stock Exchange

By R Tully for MPhil at the University of Cape Town, 2014

Cluster analysis is becoming an increasingly popular method in modern finance because of its ability to summarise large amounts of data and so help individual and institutional investors to make informed investment decisions. This study analyses the application of cluster-analysis techniques to time-series stock return data from the Johannesburg Stock Exchange. The twofold investigation first applied Salvador and Chan's L-method stopping rule to a hierarchical clustering of the data set in order to determine the number of clusters on the exchange. This implicitly confirmed the existence of a latent structure on the exchange, and established cluster analysis as a valid technique for application to Johannesburg Stock Exchange data. The second part of the analysis looked at the ability of three different clustering algorithms to generate consistent clusters and cluster members over time. The variation of information was used to measure the consistency of cluster members through time. Hierarchical clustering using Ward's method and the Euclidean distance measure proved to produce the most consistent results, while the K-means algorithms generated the least consistent cluster members.

A survey and implementation of some calibration algorithms for the SABR and Heston models

By C Webber for MPhil at the University of Cape Town, 2014

This dissertation surveys and implements some calibration methods for the SABR and Heston models. Hagan examined the effect of the SABR parameters on the skew in order to determine which parameters may be redundant. Hagan and West found that by fixing one of the parameters in the SABR model, the remaining parameters were stable over time. The author implements a SABR calibrator to confirm that the parameters are stable over time. The author then examines the effects of the five Heston parameters on the skew in order to determine if any of the parameters are redundant. Calibrators where some parameters have been fixed and calibrators where no parameters have been fixed are implemented. The performance of these calibrators is then compared based on three criteria: the stability of the parameters over time, the fit of the solution and the computational efficiency of the calibrator. The author finds that if the redundant parameters are fixed, the Heston parameters are more stable, the computation time is less and the fit is slightly worse. All implementations are done in the context of the South African market. The calibrators are programmed in Matlab and the code is included in the appendix.

Investigation of factor rotation routines in principal-component analysis of stock returns

By N Weimar for MPhil at the University of Cape Town, 2014

This paper investigates rotation routines that will produce uncorrelated rotated principal components for a dataset of stock returns, in an attempt to identify the macroeconomic factors that best explain the variability among risk-adjusted stock returns on the Johannesburg Stock Exchange. An alternative to the more traditional rotation approaches is used, which creates subsets of principal components with similar variances that are rotated in turn. It is found that only one of the three normalisation constraints examined can retain uncorrelated principal components after rotation. The results also show that when subspaces of components are rotated that have close eigenvalues, the different rotation criteria used to rotate principal components will produce similar results. After rotating the suitable subsets using varimax rotation, it is found that the first rotated component can be explained by the African industrials sector, the second rotated component is related to the African consumer services sector while the third rotated component shows a significant relationship to the African finance factor.

The effect of liquidity on stock returns on the JSE

By A Reisinger for MCom at the University of Stellenbosch, 2013

This thesis examines the effect of liquidity on excess stock returns on the Johannesburg Stock Exchange (JSE) over the period 2003 to 2011. It builds on the findings of previous studies that found size, value and momentum effects to be significant in explaining market anomalies by adding a further explanatory factor, namely liquidity. A standard CAPM, as well as a momentum-augmented-Fama–French model are employed to

perform regression analyses to examine the effect of the four variables on excess stock returns. Results suggested that the log of the stock's market value best captured the size effect, the earnings yield best captured the value effect and the previous three months' returns best captured the momentum effect. Five liquidity proxies are used: the bid-ask spread first proposed by Amihud, turnover, the price-impact measure of Amihud and two zero-return measures proposed by Lesmond et al. Despite prior studies having found liquidity to be an influential factor, this thesis found the opposite to be true. This finding remains robust, irrespective of the type of liquidity measure used. While size, value and momentum are found to be significant to a certain extent in explaining excess stock returns over the period, liquidity is not found to be significant. This is a surprising result, given that the JSE is seen as an emerging market, which is generally regarded as illiquid. This fact is exacerbated by the fact that the JSE is a highly concentrated and therefore skewed market that is dominated by only a handful of shares. Hence, liquidity is expected to be of utmost importance. The result that liquidity is however not a priced factor on this market is therefore an important finding that requires further analysis to determine why this is the case. In addition, significant non-zero intercepts remained, indicating continued missing risk factors.

Mortality models: comparison and application in old-age populations of selected countries

By B Hu for MSc at the University of the Witwatersrand, 2014

This research examined which of the five well-known chosen extrapolative mortality models best captured the trends in old-age population mortality for different age groupings in four different countries. Mortality rates from the Human Mortality Database for the United Kingdom, Poland, Japan and Taiwan were used, encompassing males and females in the 65–89 age group. This allowed assessments to be made across developed and emerging economies, and across Europe and Asia. Comparisons were made between models to understand why some work better for some age groupings in some countries. The research considered the goodness-of-fit of these well-known mortality models to historical population mortality rates, assessed the range of projected future mortality rates, and evaluated the financial impact of mortality uncertainty on annuity prices across the subject populations.

Some of the findings which emerged were that the Booth–Maindonald–Smith model tended to work best for most of the selected populations, particularly for female and Asian populations. Perhaps surprisingly, retiring females in the emerging economies can be expected to possibly outlive males in the developed economies selected. In a low yield environment, uncertainty around mortality has a noticeable impact on the range of pricing of annuities. The extent of mortality uncertainty is expected to be less for developed than in emerging economies, and less for females than males.

Modelling for the optimal product to offer a financial services customer

By J Mukomberanwa for MSc at the University of the Witwatersrand, 2014

This study illustrates how various statistical classification models can be compared and utilised to resolve cross-selling problems encountered in a financial-services environment. Various statistical classification algorithms were deployed to model for the appropriate product to sell to a financial-services customer under a multi-classifier setting. Four models were used; namely, multinomial logistic regression, multinomial bagging with logistic regression, multinomial random forests with decision trees and error-correcting output coding. The models were compared in terms of predictive accuracy, generalisation, interpretability, ability to handle rare instances and ease of use. A weighted score for each model was obtained based on the evaluation criteria stated above and, consequently, an overall model ranking.

In terms of the data, banked customers who only had a transactional account at the start of the observation period were used for modelling process. Varying samples of the customers were obtained from different time points with the preceding six to twelve months information being used to derive the predictor variables and the following six months used to monitor product take-up.

Error-correcting output coding performed the best in terms of predictive accuracy but did not perform as well on other metrics. Overall, multinomial bagging with logistic regression proved to be the best model. All the models struggled with modelling for the rare classes. Weighted classification was deployed to improve the rare-class prediction accuracy. Classification accuracy showed significant limitation under the multi-classifier setting as it tended to be biased towards the majority class. The measure of area under the receiver operating characteristic curve as proposed by Hand and Till proved to be a powerful metric for model evaluation.