

The Ross Recovery Theorem with a Regularised Multivariate Markov Chain

Vaughan van Appel^{1,2} and Eben Maré ^{*3}

¹Department of Statistics, University of Johannesburg, P.O. Box 524, Auckland Park, 2006, South Africa, E-mail: vvanappel@gmail.com

²Department of Actuarial Science, University of Pretoria, Private Bag X20, Hatfield 0028, South Africa

³Department of Mathematics and Applied Mathematics, University of Pretoria, Private Bag X20, Hatfield 0028, South Africa, E-mail: Eben.Mare@up.ac.za

November 4, 2018

Abstract

Recently, Ross [14] derived a theorem, namely the “Recovery Theorem”, that allows for the recovery of the pricing kernel and real-world asset distribution, under particular assumptions, from a forward-looking risk neutral distribution. However, recovering the real-world distribution involves solving two ill-posed problems. In this paper, we introduce and test the accuracy of a regularised multivariate mixture distribution to recover the real-world distribution. In addition, we show that this method improves the estimation accuracy of the real-world distribution. Furthermore, we carry out an empirical study, using weekly South African Top40 option trade data, to show that the recovered distribution is in line with economic theory.

Key words: Ross recovery theorem, Real-world probabilities, Regularisation, Multivariate Markov chain with a regularisation term;

1 Introduction

Asset distributions are vitally important to solve financial problems in risk management, portfolio optimisation and optimal trading strategies. A commonly used approach to forecast returns is to use historical data or opinion polling to estimate asset distributions. However, financial markets are quite volatile, and using historical distributions for forecasting are not always desirable. An alternative forecasting method is to extract the forward-looking risk-neutral distribution from the option market data. It is well known that option prices convey some market risk forecast as payoffs extend out in time. Therefore, option prices are, by nature, forward-looking. In a complete market, Black and Scholes [4] and Merton [13] proved that the value of an option is independent of the expected return on the underlying asset. This gave rise to the risk-neutral valuation framework, where the only unknown parameter affecting the option price is the assets’ underlying volatility, or commonly referred to as

*Correspondence: Eben.Mare@up.ac.za

the implied volatility. Furthermore, Breeden and Litzenberger [6] showed that the forward-looking risk-neutral distribution can be derived by option prices under the assumption of complete markets. However, the risk-neutral distribution mostly differs from the real-world distribution, which expresses market participants' consensus. In short, under the risk-neutral measure, the expected return of the asset is the risk-free rate, since the risk-neutral measure is the real measure with the risk premium removed.

While financial institutions have long used implied volatilities to gauge the market's perception of risk, option prices have been elusive to predict the real-world distribution. Recently, Ross [14] published a remarkable theorem that recovers the real-world probability distribution and pricing kernel from option prices under a particular set of assumptions. For example, one of the assumptions is that markets are complete. This is rarely true in any exchange traded option dataset, especially in South Africa, where the option price data is sparse and noisy. To satisfy this assumption, it is necessary to extrapolate forward-looking option price data (see, *e.g.*, [1, 2, 12, 8]). More specifically, Flint and Maré [8] used a deterministic SVI volatility model with a robust fitting algorithm to estimate volatility surfaces, which proved to be a promising method to estimate the forward-looking risk-neutral distribution.

It is well known that the risk-neutral probability measure is extensively used in derivative pricing, however, knowledge of the pricing kernel and real-world distribution will be invaluable for investors regrading risk management, portfolio optimisation and investment strategies. In short, the recovery theorem differs from other approaches in that it adjusts the risk-neutral distribution to a real-world distribution and does not rely on historical returns.

The empirical problem with the recovery theorem is that it is difficult to recover an accurate real-world distribution (see, *e.g.*, [2, 3, 16]), as it involves solving two ill-posed problems. The first ill-posed problem involves finding the risk-neutral distribution by taking the second derivative of the option pricing function and the second involves calculating the transition matrix that captures the dynamics of the state prices. In this paper, we will focus on the second ill-posed problem by implementing a regularised multivariate Markov chain in an attempt to stabilise the estimation of the real-world transition distribution matrix. In addition, we will conduct a numerical analysis and a robustness check to show the effectiveness of this method. Thereafter, we will apply the recovery theorem to weekly Top40 option trade data, traded on the South African Futures Exchange (SAFEX), to estimate the real-world distribution. In addition, we compare the first four moments of the real-world distribution to the risk-neutral distribution.

2 The Recovery Theorem

In this section, we start by reviewing the recovery theorem [14]. For simplicity, we adopt some of the notation and terminology used in Ross [14]. Intuitively, he attempts to recover the real-world transition probabilities of a Markovian state variable S that determines aggregate consumption, using market derivative prices on S . The derivation of the recovery theorem requires the following assumptions: (i) Transition state prices, $p_{i,j}$, need to be strictly positive, (ii) the transition state prices follow a time homogeneous process, and (iii) the corresponding pricing kernel, $\psi_{i,j}$, is transition independent.

Consider a time-homogeneous process $\{S_t\}_{t \geq 0}$ on a finite state space with n states. Since, calendar time is irrelevant, the transition probability moving from state i at time t to state j

at time $t + 1$ is given by:

$$P_{i,j} = P(S_{t+1} = j | S_t = i), \quad \forall t = 1, \dots, m - 1, \quad (2.1)$$

where P denotes a $n \times n$, one period ahead, irreducible transition matrix. If the rows of P sum to one, then we say P is a stochastic matrix; however, for the recovery theorem, P is sub-stochastic as it captures the dynamics of the discounted risk-neutral distribution, *i.e.*, state prices. Therefore, the elements, $p_{i,j}$, of the transition matrix denote the value of an Arrow-Debreu security contract that pays one unit of the numeraire if a particular state is reached in the next time step and zero otherwise. But, by normalising the rows of P to sum to unity, we define a $n \times n$ transition risk-neutral probability matrix Q , with elements:

$$q_{i,j} = \frac{p_{i,j}}{\sum_{k=1}^n p_{i,k}}, \quad \forall i, j = 1, 2, \dots, n. \quad (2.2)$$

The transition kernel, ψ , in Ross's framework is defined as the ratio price per unit of probability, *i.e.*,

$$\psi_{i,j} = \frac{p_{i,j}}{f_{i,j}}, \quad (2.3)$$

where $f_{i,j}$ is the real-world probabilities. Intuitively, one needs to solve two unknown quantities in (2.3) in order to recover the real-world probabilities. In order to do this, Ross [14] assumes that the kernel is transition independent. This assumption allows us to write the pricing kernel as

$$\psi_{ij} = \delta \frac{h(S_j)}{h(S_i)}, \quad (2.4)$$

where h is a positive function of states and δ a positive discount factor. Substituting (2.3) in (2.4) yields

$$p_{ij} = \delta \frac{h(S_j)}{h(S_i)} f_{ij}. \quad (2.5)$$

Rewriting the state equations (2.5) in matrix form, we have

$$P = \delta D^{-1} F D, \quad (2.6)$$

where P is the $n \times n$ transition probability matrix, F is the $n \times n$ real-world transition matrix, and D is the $n \times n$ diagonal matrix with the undiscounted kernel, *i.e.*,

$$D = \text{diag}(h(S_1), h(S_2), \dots, h(S_n)). \quad (2.7)$$

Solving for F in (2.6) yields

$$F = \frac{1}{\delta} D P D^{-1}. \quad (2.8)$$

Since F is a matrix whose rows are transition probabilities, *i.e.*, a stochastic matrix, we have $F \mathbf{1} = \mathbf{1}$, where $\mathbf{1}$ is a vector of ones. Using this condition, with (2.8), we have

$$P D^{-1} \mathbf{1} = \delta D^{-1} \mathbf{1}. \quad (2.9)$$

If we define the vector $\mathbf{z} \equiv D^{-1} \mathbf{1}$, we obtain

$$P \mathbf{z} = \delta \mathbf{z}. \quad (2.10)$$

If one assumes no arbitrage, then P is a non-negative matrix. Note that, if P is a positive matrix, then by definition, P is irreducible. However, if P is non-negative and all states are attainable from all other states in k steps, then P is also irreducible. Then from the Perron-Frobenius theorem there exists a unique positive eigenvector \mathbf{z} and an associated maximum eigenvalue δ . Intuitively, Ross [14] solves all three unknowns in (2.6) using the Perron-Frobenius Theorem. The following theorem guarantees a unique solution of this problem.

Theorem 2.1 (Recovery Theorem, Ross [14]). *Assuming no arbitrage, irreducibility of the pricing matrix P , and that the pricing matrix is generated by a transition independent kernel, then given any set of state prices there exists a unique positive solution pair: the pricing kernel and real-world measure.*

In short, the recovery theorem allows us to uniquely find F from P . Knowledge of the real-world distribution will be of great benefit to financial practitioners. Although, many of the assumptions in the recovery theorem are violated in real life, Audrino *et al.* [2] and Flint and Maré [8] showed by empirical studies that the real-world distribution obtained from the recovery theorem added economic value.

3 Implementation of the Ross Recovery Theorem

In this section, we describe the three step procedure, outlined in Spears [16], for implementing the recovery theorem.

Step 1: Use the method proposed by Breeden and Litzenberger [6] to construct a $n \times m$ state price matrix, S , by taking the second derivative with respect to the strike of a European call option at each tenor, *i.e.*,

$$S(K, t) = \frac{\partial^2 c(K, t)}{\partial K^2}, \quad (3.1)$$

where $c(K, t)$ is the current price of an European call option with strike, K , and tenor, t . Numerically approximating (3.1) yields the forward-looking state price function. In reality, a continuum of traded strikes is not directly observed in the markets. This is the first ill-posed problem. However, a wide range of state price estimation techniques can be found in the literature (see, *e.g.*, [1, 8, 12]). More specifically, Flint and Maré [8] used the stochastic volatility inspired (SVI) Model to model the implied volatility surface, and thus, the state price surface. Furthermore, they showed that the deterministic SVI model is a promising candidate for modelling implied volatility surfaces and ultimately estimating the underlying risk-neutral distribution. The SVI model was first introduced by Gatheral [9] and is given by

$$\sigma^2(x, t) = a + b \left(\rho(x - m) + \sqrt{(x - m)^2 + s^2} \right), \quad (3.2)$$

where $x = \ln\left(\frac{K}{F}\right)$ is the log-forward-moneyness, and the coefficients a , b , ρ , s , and m depend on the expiration and have an intuitive geometric interpretation. Furthermore, the parametrisation of the SVI model makes it relative easy to eliminate calendar spread arbitrage, making the SVI model desirable [10]. In Figure 3.1, we display an example of the implied volatility surface obtained by using the SVI model, where we can see that the SVI model provides a good interpolation of implied volatility.

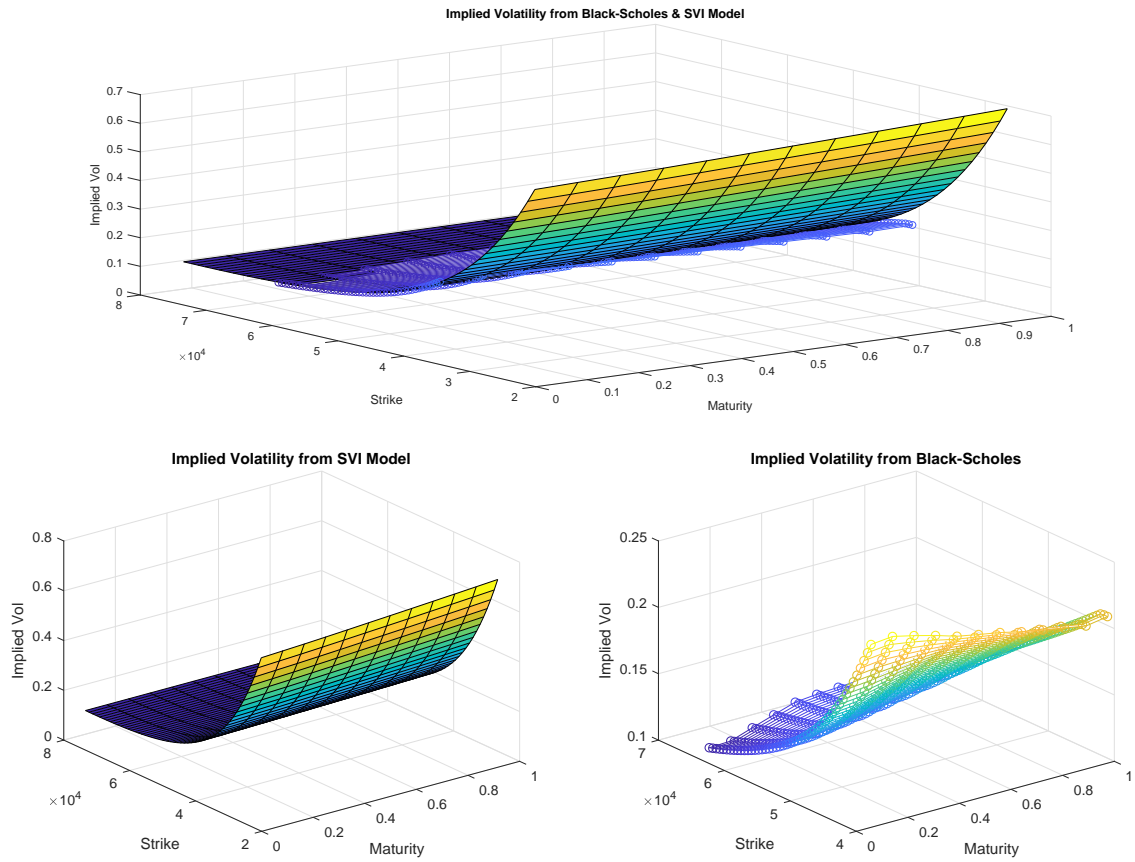


Figure 3.1: Implied volatility for the South African Top40 index on 15 January 2018: The mesh (bottom right) represent the quoted implied volatilities across maturity and strikes and the surface (bottom left) represents the implied volatilities across maturities and strikes using the SVI model. The top figure represents the overlay of the quoted and fitted implied volatilities.

After the implied volatility skews are calibrated, we can calculate the call option prices, using the Black-Scholes formula, across the full strike range for each term of the extrapolated implied volatility skews. Thereafter, using (3.1) we estimate the forward state price matrix. In Figure 3.2, we give an example of the forward state price matrix, using the extrapolated implied volatilities.




Figure 3.2: State Prices for the South African Top40 index on 15 January 2018.

Step 2: Construct a $n \times n$ state transition probability matrix, P . Unfortunately, P is not directly observed, since a rich forward market for options does not exist. However, Ross [14] shows that if $m \geq n$, we can estimate P , since it specifies a time-homogeneous transition from one maturity to the next, as follows:

$$S_{:,t+1}^\top = S_{:,t}^\top P, \quad t = 1, 2, \dots, m - 1. \quad (3.3)$$

If we denote $A = S_{:,1:m-1}^\top$ and $B = S_{:,2:m}^\top$, then (3.3) can be rewritten as an ordinary least squares (OLS) problem, as follows:

$$P = \arg \min_P \|AP - B\|_2^2 \quad (3.4)$$

$$\text{subject to } s_{:,1} = p_{i_0,:} \quad (3.5)$$

$$p_{i,j} \geq 0 \quad (i, j = 1, \dots, n), \quad (3.6)$$

where $\|\cdot\|_2$ denotes the Euclidean norm. Since $S_{:,1}$ is the one period ahead state price and P is a one period state transition matrix, we have by definition a constraint (3.5), where i_0 is the current state (normally defined at the centre of the transition matrix P , *i.e.*, $i_0 = (n + 1)/2$). In theory, equation (3.4) can easily be solved with standard optimisation techniques. Therefore, we numerically implement the OLS problem to derive the transition pricing matrix P .

Step 3: Using the Perron-Frobenius theorem, *i.e.*, (2.10), we can extract a unique positive eigenvector, \mathbf{z} , and eigenvalue, δ . Thereafter, the elements of F can be calculated using (2.8).

The accuracy of the estimation of the real-world distribution, using the recovery theorem, largely depends on how accurately the transition matrix, P , is estimated. In the literature, it has proven to be difficult to accurately estimate (3.4) and furthermore to replicate the results indicated in Ross [14]. The reason for this, is that it involves solving the second ill-posed problem, where A is ill-conditioned (*i.e.*, a small change in one of the coefficient values in A , results in a large relative change in the solution values), which renders active-set optimisation methods that are dependent on $A^\top A$ infeasible, as in this case. This can be seen in Audrino *et al.* [2], Kiriu and Hibiki [11], and Spears [16] suggesting that Ross [14] placed significant

constraints on the structure of the transition matrix. In an attempt to replicate the results in Ross [14], Spears [16] implemented nine optimisation methods for solving (3.4). Furthermore, Sanford [15] proposed a mixture transition distribution, where the proposed states depend on the current state price and its option implied volatilities to stabilise the estimation of P . More specifically, Sanford [15] simplifies the original specification of the multivariate model by assuming that contingent state prices are solely defined by the state levels, but conditioned on the volatility. That is,

$$S_{:,t+1}^\top = S_{:,t}^\top P + \sigma_t^{(IV)} \beta, \quad t = 1, 2, \dots, m-1, \quad (3.7)$$

where $\sigma_t^{(IV)}$ is the implied volatility state at time t as it is the best representation of the market's future volatility state and β is the volatility transition matrix. Furthermore, Sanford [15] shows that the multivariate method had a significant improvement on the univariate recovery theorem as the volatility acts as a proxy for economical uncertainty. Similarly, equation (3.7) can be reduced to the following general optimisation problem:

$$P = \arg \min_{P, \beta} \left\| AP + \sigma_t^{(IV)} \beta - B \right\|_2^2, \quad \forall t = 1, 2, \dots, m-1 \quad (3.8)$$

$$\text{subject to (3.5), (3.6) and } \beta \geq 0. \quad (3.9)$$

In theory, the multivariate model gives a third dimension in the Markov chain. Intuitively, more variables could be added to the regression model. However, this will come at a computational cost and the more variables added to the regression equation, will result in too few degrees of freedom to consider the resulting state price matrix, P , reliable.

An alternative method of stabilising the estimation of P is by adding a regularisation parameter to the estimation process. This has proven to be a successful method in the studies conducted by Audrino *et al.* [2] and Kiriu and Hibiki [11]. Therefore, this paper contributes in two ways. Firstly, we compare the multivariate method with the regularised methods (this has not been done to our knowledge) and secondly, due to the success of the regularised methods in the literature, we extend the multivariate method by adding a regularisation term.

3.1 Ridge Regularisation Methods

An effective method in stabilising the estimation of the transition matrix, P , is to introduce a regularisation term. The use of a regularisation term to solve ill-posed problems was first introduced by Tikhonov [17]. The Tikhonov method is a standard regularisation method used in the literature to solve ill-posed problems.

3.1.1 Tikhonov Regularisation without Prior Information

In this section, we review two regularisation methods to estimate P , found in the literature, and extend the multivariate method by adding a regularisation term. Audrino *et al.* [2] first introduced the implementation of the Tikhonov regularisation (ridge regression) method in estimating P in the recovery theorem, by the following constrained optimisation problem:

$$P = \arg \min_P \|AP - B\|_2^2 + \zeta \|P\|_2^2 \quad (3.10)$$

$$\text{subject to (3.5) and (3.6),} \quad (3.11)$$

where ζ is a regularisation parameter that controls the trade-off between fitting and stability. The selection method of ζ is paramount in finding an accurate solution. Therefore, Audrino *et al.* [2] proposed that an optimal ζ can be determined by minimising the discrepancy between the observable state price matrix (S^O) and the unrolled state price matrix (S^P) implied by matrix P , *i.e.*,

$$S_{:,t}^{P^\top} = \iota_{i_0}^\top P^t, \quad t = 1, 2, \dots, m, \quad (3.12)$$

where ι_{i_0} denotes a vector with 1 in the i_0^{th} position and zeros elsewhere, and P^t denotes the t -steps ahead state approximation. Furthermore, they use the Kullback-Lieber (KL) divergency as a measure of discrepancy between the two matrices, by solving ζ that minimises:

$$\arg \min_{\zeta} D_{KL}(S^O || S^P), \quad (3.13)$$

where

$$D_{KL}(S^O || S^P) = \sum_{i=1}^n \sum_{t=1}^m S_{i,t}^O \ln \left(\frac{S_{i,t}^O}{S_{i,t}^P} \right) - \sum_{i=1}^n \sum_{t=1}^m S_{i,t}^O + \sum_{i=1}^n \sum_{t=1}^m S_{i,t}^P, \quad (3.14)$$

and the optimal ζ is derived iteratively.

Note that equation (3.10) can be we rewritten as a constraint OLS problem as follows [2]:

$$P = \arg \min_{P \geq 0} \left\| \begin{bmatrix} A \\ \sqrt{\zeta} I \end{bmatrix} P - \begin{bmatrix} B \\ O \end{bmatrix} \right\|_2^2, \quad (3.15)$$

where I denotes an identity matrix and O is a vector of zeros. In an empirical study using daily closing prices of out-of-the-money call and put options on the S&P 500 for each Wednesday between 5 January 2000 and 26 December 2012, Audrino *et al.* [2] showed that the Thikonov regularisation drastically improved the stability of the estimation of the transition matrix and showed that there is economic value in the recovered distributions.

In the next section, Kiriu and Hibiki [11] extended the estimation of P by using the Tikhonov regularisation method with prior information.

3.1.2 Tikhonov Regularisation with Prior Information

The second regularisation method we review in this study was introduced by Kiriu and Hibiki [11], where they extended the regularisation term above to consider prior information. For the prior information, \bar{P} , they suggest that $p_{i,j}$ should be similar to $p_{i+k,j+k}$ for all $k \leq \min(n-i, n-j)$. Furthermore, they estimated P , using a problem specific error function in an attempt to balance the relative gain in the objective function from each term in the regularised optimisation problem, as follows:

$$P = \arg \min_{P \geq 0} \|AP - B\|_2^2 + \zeta \|P - \bar{P}\|_2^2 \quad (3.16)$$

$$= \arg \min_{P \geq 0} y^{\text{fit}}(\zeta) + \zeta y^{\text{reg}}(\zeta) \quad (3.17)$$

$$\text{subject to (3.5) and (3.6),} \quad (3.18)$$

where

$$\bar{P} = \begin{bmatrix} \sum_{k=1}^{i_0} s_{k,1} & s_{i_0+1,1} & \cdots & s_{n-1,1} & s_{n,1} & 0 & \cdots & 0 & 0 \\ \vdots & \vdots & \cdots & \vdots & \vdots & \vdots & \cdots & \vdots & \vdots \\ \sum_{k=1}^2 s_{k,1} & s_{3,1} & \cdots & s_{i_0,1} & s_{i_0+1,1} & s_{i_0+2,1} & \cdots & s_{n,1} & 0 \\ s_{1,1} & s_{2,1} & \cdots & s_{i_0-1,1} & s_{i_0,1} & s_{i_0+1,1} & \cdots & s_{n-1,1} & s_{n,1} \\ 0 & s_{1,1} & \cdots & s_{i_0-2,1} & s_{i_0-1,1} & s_{i_0,1} & \cdots & s_{n-2,1} & \sum_{k=n-1}^n s_{k,1} \\ \vdots & \vdots & \cdots & \vdots & \vdots & \vdots & \cdots & \vdots & \vdots \\ 0 & 0 & \cdots & 0 & s_{1,1} & s_{2,1} & \cdots & s_{i_0-1,1} & \sum_{k=i_0}^n s_{k,1} \end{bmatrix}, \quad (3.19)$$

$y^{\text{fit}}(\zeta)$ represents the fitting error and $y^{\text{reg}}(\zeta)$ represents the deviation between P and \bar{P} . Furthermore, Kiriu and Hibiki [11] showed that as ζ increases, y^{fit} decreases and y^{reg} increases monotonically. Therefore, they selected ζ by minimising the problem specific function:

$$h(\zeta) = \frac{y^{\text{fit}}(\zeta) - y^{\text{fit}}(0)}{y^{\text{fit}}(\infty) - y^{\text{fit}}(0)} + \frac{y^{\text{reg}}(\zeta) - y^{\text{reg}}(\infty)}{y^{\text{reg}}(0) - y^{\text{reg}}(\infty)}, \quad (3.20)$$

where the denominators represents the maximum spread in each term and the numerator represents the spread for a specified ζ value.

In addition, Kiriu and Hibiki [11] compared the effectiveness of this selection method with (3.14), where they found that (3.20) yielded better results. Therefore, for the remainder of this study, we will use (3.20) as the selection method for ζ . Equation (3.16) can also be formulated as an OLS problem, as follows [11]:

$$P = \arg \min_{P \geq 0} \left\| \begin{bmatrix} A \\ \sqrt{\zeta} I \end{bmatrix} P - \begin{bmatrix} B \\ \sqrt{\zeta} \bar{P} \end{bmatrix} \right\|_2^2. \quad (3.21)$$

In a simulated study, Kiriu and Hibiki [11] showed that their method estimated the real-world distribution more accurately than the Tikhonov method proposed by Audrino *et al.* [2]. Furthermore, in a similar empirical study to Audrino *et al.* [2], Flint and Maré [8] implemented the regularisation method with prior information to extract the real-world distribution on a history of implied volatility surfaces for the South African Top40 index, where they showed that the recovered real-world moments are in line with economic rationale and showed promising results when used in a simple asset allocation framework.

Since the regularisation methods have proven to be a powerful method in estimating the real-world distribution in the recovery theorem, we extend the multivariate method to a regularised multivariate method in the next section.

3.1.3 The Multivariate Model with a Tikhonov Regularisation

Later, we will show that the addition of the regularisation term in the estimation procedure improves the estimation of P and ultimately F (see also, *e.g.*, [2, 11]). Therefore, we extend the multivariate Markov chain proposed by Sanford [15] to a regularised multivariate Markov

chain by adding the regularisation parameter as follows:

$$P = \arg \min_{P, \beta} \left\| AP + \sigma_t^{(IV)} \beta - B \right\|_2^2 + \zeta \|P\|_2^2, \quad \forall t = 1, 2, \dots, m-1 \quad (3.22)$$

$$\text{subject to (3.5), (3.6) and } \beta \geq 0. \quad (3.23)$$

Furthermore, we also extend the optimisation problem above, with the regularisation of prior information, as such,

$$P = \arg \min_{P, \beta} \left\| AP + \sigma_t^{(IV)} \beta - B \right\|_2^2 + \zeta \|P - \bar{P}\|_2^2, \quad \forall t = 1, 2, \dots, m-1 \quad (3.24)$$

$$\text{subject to (3.5), (3.6) and } \beta \geq 0, \quad (3.25)$$

where \bar{P} is given in (3.19). We found that the regularised method with prior information performed better than the regularised method without prior information. Therefore, we will only consider (3.24) in the remainder of this paper.

3.2 Elastic Net Regularisation Method

Elastic net is a regression regularisation method used in statistics, that linearly combines the L_1 and L_2 penalties of the lasso and ridge methods. The (L_1) penalty achieves sparsity in the model by setting the irrelevant regression coefficient equal to zero and the (L_2) penalty achieves robustness in the model. Therefore the optimisation problem becomes:

$$P = \arg \min_{P \geq 0} \|AP - B\|_2^2 + \zeta \|P\|_2^2 + \lambda \|P\|_1 \quad (3.26)$$

$$\text{subject to (3.5) and (3.6),} \quad (3.27)$$

where the estimation is carried out in a two-stage procedure as follows: for each fixed ζ , it finds the ridge regression coefficients and then does a lasso shrinkage along the lasso coefficient path [18]. Furthermore, Zou and Hastie [18] refer to this as the naïve elastic net criterion, since it appears to amount to double shrinkage, where it was found that the naïve elastic net regularisation method does not perform well, unless it is close to ridge or lasso. In our study, we find that λ is small indicating that it is close to ridge. However, to improve the prediction performance, Zou and Hastie [18] rescale the coefficients of the naïve version of elastic net by multiplying the estimated coefficients by $(1 + \lambda_2)$. Next, we add the prior information to (3.26), yielding

$$P = \arg \min_{P \geq 0} \|AP - B\|_2^2 + \zeta \|P - \bar{P}\|_2^2 + \lambda \|P\|_1 \quad (3.28)$$

$$\text{subject to (3.5) and (3.6).} \quad (3.29)$$

The elastic net with prior information yielded better results than without the prior information. Therefore, for the remainder of this study, we will only show the results for the elastic nets with prior information.

In the next section, we compare the estimation methods discussed above by estimating the real-world distribution, where we will show that the regularised multivariate method gives a better estimate than the methods reviewed by conducting a similar simulation study to Kiriu and Hibiki [11].

4 Comparison of Methods

In this section, we compare the accuracy of the estimation of P , using the methods discussed in Section 3. We will follow the same estimation accuracy procedure and robust check outlined in Kiriu and Hibiki [11] as follows:

1. Firstly, a hypothetical real-world matrix (F^H) is obtained from the historical daily S&P 500 index price data. More specifically, we set 11 returns (states) in total, placed every 6% symmetrically around 0%. F^H is generated by setting a reference date and calculating 12 returns every 30 calendar days, where the S&P 500 returns are calculated as follows:

$$\text{Return} = \log \left(\frac{S_T}{S_0} \right) \cdot 100\%.$$

A matrix is generated by calculating the number of state transitions of the return in one period. This is repeated daily by changing the reference date from 02 January 1986 to 30 December 2016. Thereafter, all matrices are summed up and divided by the summed matrix row total, giving an 11×11 probability matrix (see Figure 4.1).




Figure 4.1: F^H

Secondly, the pricing kernel matrix (Φ^H) is obtained by assuming that the investor has a CRRA utility function, $U(c) = c^{1-\gamma_R}/(1-\gamma_R)$, with relative risk aversion γ_R , *i.e.*,

$$\phi_{i,j} = \delta \left(\frac{1+r_j}{1+r_i} \right)^{-\gamma_R} \quad \forall i, j = 1, \dots, n, \quad (4.1)$$

where $\gamma_R = 3$ and $\delta = 0.999$. These parameters were chosen to be consistent with the parameters reported in Bliss and Panigirtzoglou [5], where they estimated the risk aversion parameter, γ , implied in the S&P 500 option data and historical option price data from 1993 to 2010, to have a minimum risk aversion parameter value to be 3.37 and a maximum value of 9.52. The maximum parameter value will be used in the robust check in the Section 4.2.

2. A hypothetical transition state price matrix P^H (see Figure 4.2) is calculated backward from the matrices F^H and Φ^H , *i.e.*,

$$P^H = \Phi^H F^H. \quad (4.2)$$




Figure 4.2: P^H

3. A hypothetical current state price matrix S^H (see Figure 4.3) is calculated backward from the matrix P^H , *i.e.*,

$$S_{:,j+1}^H = (S_{:,j}^H)^\top P^H, \quad \forall j = 1, \dots, m, \quad (4.3)$$

where $S_{:,1}^H = P_{i_0, :}^H$.




Figure 4.3: S^H

4. White noise is added to S^H to obtain S^N , as follows:

$$S_{i,j}^N = S_{i,j}^H(1 + e_{i,j}), \quad \forall i, j = 1, 2, \dots, n, \quad (4.4)$$

where $e_{i,j} \sim N(0, \sigma)$.

5. Estimate P^N from S^N , using (3.4), (3.8), (3.15), (3.21), (3.24) and (3.28). In the case of the multivariate estimation methods, we will use a flat implied volatility, $\sigma^{(IV)}$, of 10%. We note that more accurate results could be achieved by modelling the behaviour of volatility and incorporating a forward-looking volatility structure than only looking at a flat or current volatility.

6. F^N is derived by applying the recovery theorem for each of the estimated matrices P^N .
7. The closer the estimated real-world distribution matrix F^N is to F^H , the more accurate the estimation process is.

Next, in order to measure how close the two distributions are we use the Kullback-Leibler Divergence test.

4.1 Kullback-Leibler Divergence

Intuitively, we would like to measure how close we can get back to F^H using S^N . Therefore, we will follow the same estimation accuracy method outlined in Kiri and Hibiki [11], namely, the Kullback-Leibler (KL) divergence test. The KL divergence test measures the difference between two distributions and is given as follows:

$$D_{KL}(F^N|F^H) = \sum_{i=1}^n \sum_{j=1}^n f_{i,j}^N \ln \left(\frac{f_{i,j}^N}{f_{i,j}^H} \right). \quad (4.5)$$

Obviously, when the estimated distribution and true distribution are exactly the same, the D_{KL} will equal zero. In Figure 4.4, we show the log-log plots of the KL divergence at current state and at full state for five estimation methods discussed in this study for the regularisation parameter, $\zeta = 10^{-8}, 10^{-0.75}, \dots, 10^{1.75}, 10^2$ and $\sigma = 5\%$. In addition, we also show the KL divergence for the risk-neutral distribution (RND). The RND, Q , is the distribution obtained when using P^H in (2.2). Note that this is the best possible estimate for the RND as P^H is used. Therefore, obtaining a KL divergence less than the KL divergence for the RND will indicate that the estimation of the real-world distribution is more beneficial than the RND. For the current state (see Figure 4.4a), we found that both the basic method and the multivariate method provided a worse estimate of the real-world distribution than the RND. However, this is not the case for the estimation methods with the regularisation term (see Figure 4.4a). The regularisation methods clearly outperform the non-regularised methods, where the multivariate regularisation method, proposed in this paper, yielded the smallest KL divergence at current state. Similarly, the two regularisation methods with prior information clearly yield a lower KL divergence at full state compared with the basic, multivariate, and Tikhonov regularisation methods without prior information. However, the RND yielded the lowest KL divergence at full state (see Figure 4.4b).

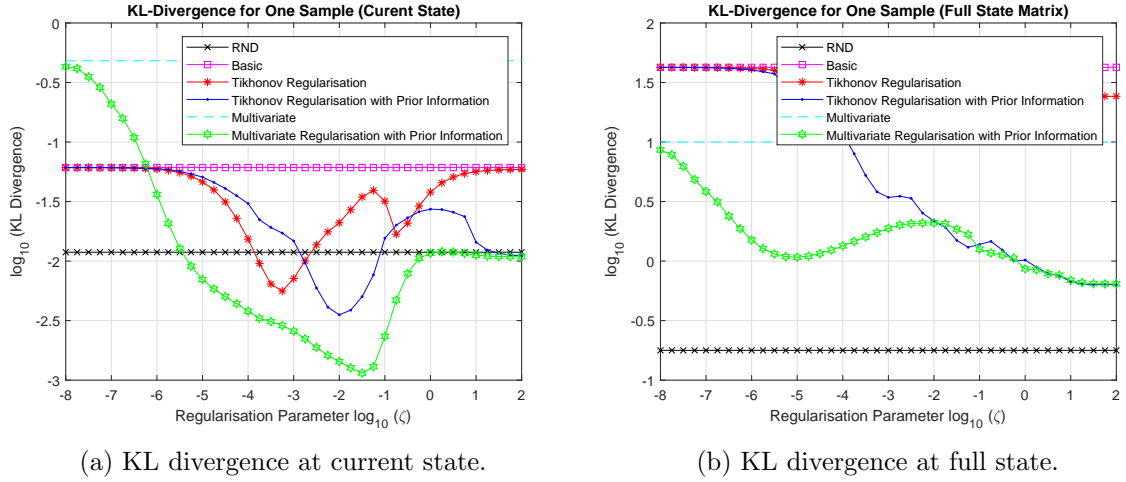


Figure 4.4: KL divergence of the real-world transition matrix.

In Figure 4.5, we show that $h(\zeta)$ is a smooth and continuous function, where a minimum value can easily be estimated, making it an appealing selection function.




Figure 4.5: Optimisation of h_k .

Next, we examine the effectiveness of the estimation methods and the selection criteria when the regularisation term is added by carrying out 1000 Monte Carlo simulations. In Table 1, we show that the expected KL divergence and standard error for the 1000 Monte Carlo simulations for the current state, *i.e.*, the i_0^{th} row vector of matrix F . More specifically, $\mathbb{E}(KL)$ represents the expected KL divergence, $\mathbb{E}(KL_{\min h_k})$ represents the expected KL divergence, where $h(\zeta)$ is a minimum, and $\mathbb{E}(\min KL)$ represents the minimum KL divergence across all ζ . We can see that the RND provides a better estimation, with a lower KL divergence, than the basic and multivariate estimation methods (as seen in Figure 4.4a). This is a direct consequence of the ill-posed problem, when solving (3.4). The regularised methods clearly outperformed the RND, basic, and multivariate methods, indicating the strength of adding the regularising term when solving ill-posed problems. More specifically, the multivariate

regularised method and the elastic net method, proposed in this paper, yielded the best results with the lowest expected KL divergence. In all cases the standard errors are small indicating the estimation methods provide stable estimates. However, it must be noted that the elastic net method is significantly more computationally expensive than the other methods discussed in this study.

Method		$\sigma = 5\%$	$\sigma = 10\%$
RND	$\mathbb{E}(KL)$	0.0119	0.0119
Basic (3.4)	$\mathbb{E}(KL)$	0.1931	0.2155
	SE	0.0067	0.0073
Multivariate (3.8)	$\mathbb{E}(KL)$	0.5908	0.7702
	SE	0.0038	0.0168
Regularised (3.15)	$\mathbb{E}(KL_{\min h_k})$	0.0335 (0.5378)	0.0344 (0.4674)
	SE	0.0002 (0.0049)	0.0004 (0.0046)
	$\mathbb{E}(\min KL)$	0.0101 (0.0256)	0.0165 (0.0640)
	SE	0.0002 (0.0020)	0.0003 (0.0028)
Regularised with Prior Information (3.21)	$\mathbb{E}(KL_{\min h_k})$	0.0124 (0.0179)	0.0494 (0.0154)
	SE	0.0003 (0.0007)	0.0017 (0.0006)
	$\mathbb{E}(\min KL)$	0.0061 (5.9265)	0.0111 (41.3819)
	SE	0.0001 (0.7453)	0.0001 (1.5536)
Multivariate Regularised with Prior Information (3.24)	$\mathbb{E}(KL_{\min h_k})$	0.0082 (0.0156)	0.0532 (0.0253)
	SE	0.0005 (0.0002)	0.0051 (0.0005)
	$\mathbb{E}(\min KL)$	0.0034 (0.0747)	0.0072 (6.7491)
	SE	0.0001 (0.0020)	0.0002 (0.7843)
Elastic Net (3.26)	$\mathbb{E}(\min KL)$	0.0031 (0.0178)	0.0062 (0.0562)
	SE	0.0001 (0.0000)	0.0001 (0.0000)

* ζ displayed in parenthesis

Table 1: KL divergence at current state matrix.

Similarly, in Table 2 we show that the expected KL divergence and standard error for the entire F matrix. We see that the multivariate method yields a smaller KL divergence than the basic and regularised method proposed by Audrino *et al.* [2]. However, the methods that are regularised with prior information still yielded the lowest KL divergence, with the multivariate regularised method yielding the lowest expected KL divergence, where $h(\zeta)$ is a minimum. Furthermore, the elastic net method yielded the lowest KL divergence across all regularisation parameters. However, the elastic net method is substantially more computationally expensive than the multivariate regularised method and, therefore, will not be studied any further in this paper.

Method		$\sigma = 5\%$	$\sigma = 10\%$
RND	$\mathbb{E}(KL)$	0.1779	0.1779
Basic (3.4)	$\mathbb{E}(KL)$	43.7850	58.3776
	SE	0.2718	0.3417
Multivariate (3.8)	$\mathbb{E}(KL)$	10.3142	10.8718
	SE	0.0136	0.0580
Regularised (3.15)	$\mathbb{E}(KL_{\min h_k})$	18.9698 (0.5378)	18.6056 (0.4674)
	SE	0.0229 (0.0049)	0.0268 (0.0046)
	$\mathbb{E}(\min KL)$	11.3012 (0.0055)	12.5754 (0.0162)
	SE	0.0270 (0.0002)	0.0308 (0.0005)
Regularised with Prior Information (3.21)	$\mathbb{E}(KL_{\min h_k})$	2.3659 (0.0179)	3.7873 (0.0154)
	SE	0.0236 (0.0007)	0.0582 (0.0006)
	$\mathbb{E}(\min KL)$	0.6828 (75.5576)	0.7064 (65.1735)
	SE	0.0010 (1.1127)	0.0022 (1.2155)
Multivariate Regularised with Prior Information (3.24)	$\mathbb{E}(KL_{\min h_k})$	2.2282 (0.0156)	3.6187 (0.0253)
	SE	0.0190 (0.0002)	0.0564 (0.0005)
	$\mathbb{E}(\min KL)$	0.6853 (87.3217)	0.7130 (87.2990)
	SE	0.0010 (0.6783)	0.0021 (0.6926)
Elastic Net (3.26)	$\mathbb{E}(\min KL)$	0.6567 (49.7875)	0.6758 (54.0514)
	SE	0.0010 (0.9809)	0.0020 (1.2007)

* ζ displayed in parenthesis

Table 2: KL divergence at Full state matrix.


It is evident from the above that the multivariate regularised method introduced in this paper improved the estimation of the real-world distribution. It must also be noted that the further the row, in the state transition matrix, is from the current state's row (*i.e.*, normally defined as the middle row), the more difficult it is to determine, but also the less influential it is on the real-world distribution (see [3]). Therefore, the transition from the current state is of greater interest in this study as we are mostly interested in how the asset would change over one period given today's state. In the next section, we conduct a robust check.

4.2 Robust Check


In this section, we conduct a robust check by using different hypothetical data obtained from the real-world distribution used above [11]. More specifically, Figures 4.6a-4.6b shows the KL divergence where $\delta = 0.995$, Figures 4.6c-4.6d shows the results for a large risk aversion parameter, namely, $\gamma = 10$ and lastly Figures 4.6e-4.6f shows the KL divergence using the CARA utility function, *i.e.*,

$$\phi_{i,j} = \delta e^{-\gamma(r_j - r_i)}, \quad i, j = 1, \dots, n \quad (4.6)$$


with $\gamma = 3$, instead of CRRA utility function. The results obtained in Figure 4.6 shows that the multivariate regularised method, proposed in this paper, yields a robust estimate of the real-world distribution. Furthermore, we carry out the robust check on the South African Top40 index, where we obtained similar results (see Figures 4.6g-4.6h).




(a) current state: $\delta = 0.995$




(b) full state: $\delta = 0.995$




(c) current state: $\gamma_R = 10$



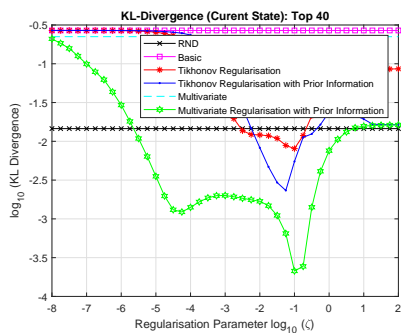
(d) full state: $\gamma_R = 10$



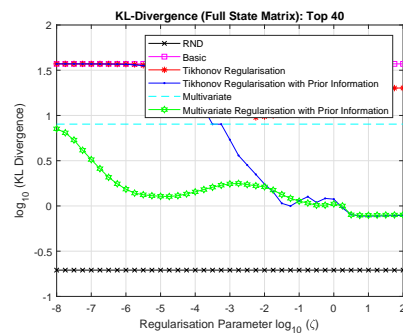
(e) current state: CARA utility



(f) full state: CARA utility



(g) current state: Top40



(h) full state: Top40

Figure 4.6: KL divergence: Robust Test.

We note that other norms, such as $\|\cdot\|_1$ and $\|\cdot\|_\infty$ could be used to estimate P more accurately. Chvátal [7] asserts that when estimating linear function, $\|\cdot\|_1$ gives the most robust answer, $\|\cdot\|_\infty$, avoids gross discrepancies with the data, and if the errors are known to be normally distributed then $\|\cdot\|_2$ is the best choice. However, in our analysis, we found that the Euclidean norm yielded the most accurate and stable results. In the next section, we conduct an empirical study.

5 Empirical Results

In this section, we compare some distributional properties of the risk-neutral and real-world distributions by using the weekly Top40 option trade data, traded on the South African Futures Exchange (SAFEX). We start by using weekly arbitrage-free implied volatility surfaces to estimate the risk-neutral distribution over the period 5 September 2005 - 15 January 2018. Furthermore, we used the SVI model to interpolate over the fixed domain $\psi \in [0.5, 1.5]$, where ψ is defined as the spot moneyness (*i.e.*, $\psi = K/S_0$) and $T \in [1, 12]$ as outlined in Flint and Maré [8]. The evolution of the weekly one-month percentiles and mean of the risk-neutral distribution is shown in Figure 5.1.






Figure 5.1: Weekly one-month percentiles of the risk-neutral Top40 distributions, 05 Sep 2005 - 15 Jan 2018.

As expected the risk-neutral distribution widened over the global financial crisis (2008-2009) and has since narrowed considerably. Next, we estimated the transition probability matrix, P , using the methods proposed by Kiriu and Hibiki [11] and the regularised multivariate method with prior information. Thereafter, we applied the recovery theorem. In Figure 5.2, we show how the risk-neutral and recovered real-world distributions widened during the financial crises (03 November 2008) compared to the risk-neutral and real-world distributions after the financial crisis (15 January 2018).




(a) 03 Nov. 2008




(b) 15 Jan. 2018

Figure 5.2: Risk-neutral and real-world distributions.


In Figure 5.3, we show the evolution over time of the weekly one-month first four moments. We see that the expected returns of the two real-world distributions are mostly above the risk-neutral distributions expected returns, except during the financial crisis. The volatility has steadily decreased since the global financial crisis (a peak of approximately 17% down to 4%). In addition, the real-world distribution obtained by using the regularised multivariate Markov chain with prior information showed a lower volatility than the distributions obtained using the univariate regularised method with prior information and risk-neutral (which showed similar volatility). This is somewhat expected, since controlling the volatility in the multivariate regression model provided us with a better sense of future economical uncertainty (see, *e.g.*, [15]). The skewness for the risk-neutral distribution became less negative during the financial crisis along with a drop in kurtosis. The skewness has since reverted to a skewness around -0.5 along with an increase in kurtosis. In addition, the weekly skewness coefficients for the real-world distributions showed sharp spikes (became positively skewed) in 2012 and 2016.




(a) Mean Return



(b) Volatility



(c) Skewness



(d) Kurtosis

In Table 3, we show the mean and volatility for the Top40 index with the first four moments of the risk-neutral and real-world distributions. The recovered moments estimated from option prices clearly provides insight above the risk-neutral moments. Furthermore, we found that the recovered kurtosis of the real-world distribution using the Tikhonov regularisation method with prior information was considerably more volatile over time than the multivariate regularisation method with prior information and the RND.

	Mean	Volatility
Top40 Returns*	28.87%	20.20%
<i>Risk-Neutral Distribution</i>		
Expected Return*	4.07%	2.20%
Volatility on Return*	22.38%	7.34%
Skewness	-0.59	0.32
Kurtosis	6.89	3.67
<i>Real-World Distribution: Tikhonov Regularisation with Prior Information</i>		
Expected Return*	13.17%	11.09%
Volatility on Return*	22.32%	6.98%
Skewness	-0.23	0.54
Kurtosis	10.49	8.08
<i>Real-World Distribution: Multivariate Regularisation with Prior Information</i>		
Expected Return*	15.05%	14.19%
Volatility on Return*	20.62%	7.18%
Skewness	-0.58	0.46
Kurtosis	4.87	3.32

*Values are annualised

Table 3: Top40 weekly one-month moments.

The predictive information obtained using the recovery theorem along with real-world data surely yielded some insight into the markets subjective probabilities. However, the true practicality and usefulness of the model remains elusive in the literature.

6 Conclusion

The recovery theorem is a remarkable theorem that allows us to estimate the real-world distribution from the risk-neutral distribution. However, the implementation of the recovery theorem requires the solution of two ill-posed problems. The first is estimating the state price matrix by calculating the second partial derivative of the option price with respect to the strike. This is especially problematic in noisy and sparse markets. Flint and Maré [8] proposed an algorithm for this first ill-posed problem. The second entails the estimation of the transition price matrix that captures the state price dynamics. Audrino *et al.* [2] and Kiriu and Hibiki [11] used a regularisation technique to obtain a stable transition matrix. In addition, Audrino *et al.* [2] and Flint and Maré [8] showed by empirical work that there is information contained in the recovered distributions. In this study, we investigated several estimation methods to accurately estimate the transition price matrix. The accuracy of the estimated transition matrix has a significant impact on the estimation of the real-world

distribution implied from option prices using the Ross recovery theorem. In addition, we presented a regularised multivariate Markov chain with prior information to estimate the transition matrix. This is a first attempt to regularize the multivariate Markov chain for the recovery theorem. In our analysis, we found that the regularised multivariate Markov chain method improved upon the estimation of the real-world distribution. Furthermore, we conducted an empirical study using weekly South African Top40 option trade data to estimate the risk-neutral and real-world distributions.

Acknowledgements

The authors would like to express their gratitude towards the editorial team as well as the two anonymous referees that aided in the finalisation of this article. The authors would also like to express their gratitude towards Emlyn Flint from Peregrine Securities for assistance with historical option price data and helpful discussions.

References

- [1] AÏT-SAHALIA Y & LO AW, 1998, *Nonparametric Estimation of State-Price Densities Implicit in Financial Asset Prices*, *The Journal of Finance*, **53**(2), pp. 499–547
- [2] AUDRINO F, HUITEMA R & LUDWIG M, 2014, *An Empirical Analysis of the Ross Recovery Theorem*, Available at SSRN: <https://ssrn.com/abstract=2433170> or <http://dx.doi.org/10.2139/ssrn.2433170>.
- [3] BACKWELL A, 2015, *State Prices and Implementation of the Recovery Theorem*, *The Journal of Risk and Financial Management*, **8**, pp. 2–16, <https://doi.org/10.3390/jrfm8010002>.
- [4] BLACK F & SCHOLES M, 1973, *The Pricing of Options and Corporate Liabilities*, *Journal of Political Economy*, **81**(3), pp. 673–654.
- [5] BLISS RR & PANIGIRTZOGLOU N, 2004, *Option-Implied Risk Aversion Estimates*, *The Journal of Finance*, *The Journal of Finance*, **59**(1), pp. 407–446.
- [6] BREEDEN DT & LITZENBERGER RH, 1978, *Prices of State-Contingent Claims Implicit in Option Prices*, *Journal of Business*, **51**(4), pp. 621–651.
- [7] CHVÁTAL V, 1983, *Linear Programming*, Freeman, New York.
- [8] FLINT E & MARÉ E, 2017, *Estimating Option-Implied Distributions in Illiquid Markets and Implementing the Ross Recovery Theorem*, *South African Actuarial Journal*, **17**(1), pp. 1–28, <http://dx.doi.org/10.4314/saaj.v17i1.1>.
- [9] GATHERAL J, 2004, *A Parsimonious Arbitrage-Free Implied Volatility Parametrization with Application to the Valuation of Volatility Derivatives*, Presentation at Global Derivatives and Risk Management, Madrid.
- [10] GATHERAL J & JACQUIER A, 2014, *Arbitrage-free SVI Volatility Surfaces*, *Quantitative Finance*, **14**(1), pp. 59–71.

- [11] KIRIU T & HIBIKI N, 2015, *Estimating Forward Looking Distribution with the Ross Recovery Theorem*, Proceedings of the Asia Pacific Industrial Engineering & Management Systems Conference.
- [12] MELICK WR & THOMAS CP, 1997, *Recovering an Asset's Implied PDF from Option Prices: An Application to Crude Oil during the Gulf Crises*, Journal of Financial and Quantitative Analysis, **32**(1), pp. 91–115.
- [13] MERTON RC, 1973, *The of Rational Option Pricing*, The Bell Journal of Economics and Management Science, **4**(1), pp. 141–183.
- [14] ROSS S, 2015, *The Recovery Theorem*, The Journal of Finance, **70**(2), pp. 615–647, <https://doi.org/10.1111/jofi.12092>.
- [15] SANFORD A, 2017, *Recovery Theorem with a Multivariate Markov Chain*, Working paper, University of Washington.
- [16] SPEARS TJ, 2013, *On Estimating the Risk-Neutral and Real-World Probability Measures*, Masters Thesis, University of Oxford.
- [17] TIKHONOV AN, 1973, *Solution of Incorrectly Formulated Problems and the Regularization Method*, Soviet Math. Dokl., **4**(4), pp. 1035–1038.
- [18] ZOU H & HASTIE T, 2015, *Regularization and Variable Selection via the Elastic Net*, Journal of the Royal Statistical Society: Series B, **67**(2), pp. 301–320.