Revealing the Implied Risk-neutral MGF from Options: the Wavelet Method

Emmanuel Haven^a, Xiaoquan Liu^b, Chenghu Ma^c, Liya Shen^b

^aUniversity of Leicester, Leicester LE1 7RH, UK*

^bUniversity of Essex, Colchester CO4 3SQ, UK

^cXiamen University, Xiamen, P. R. China

September 2, 2008

Abstract

Options are believed to contain unique information on the risk-neutral moment generating function (MGF) or the risk-neutral probability density function (PDF) of the underlying asset. This paper applies the wavelet method to approximate the implied risk-neutral MGF from option prices. Monte Carlo simulations are carried out to show how the risk-neutral MGF can be obtained using the wavelet method. With the Black-Scholes model as the benchmark, we offer a novel method to reveal the implied MGF, and to price in-sample options and forecast out-of-sample option prices with the estimated MGF.

JEL Classification: C02, C63, G13

Keywords: Wavelet analysis, option pricing, Laplace transform

*Corresponding author. E-mail: e.haven@leicester.ac.uk. Phone: +44 (0)116 2523955.

Fax: +44 (0)116 2525515

1 Introduction

In the early 1970s, Black and Scholes (1973) published the seminal Black-Scholes option pricing formula, one of the most important advances in option pricing. However, since the 1987 stock market crash there is growing empirical evidence that the market differs systematically from the Black-Scholes paradigm. There are mainly two stylized facts: (1) The Black-Scholes model assumes that the volatility of the underlying security is constant. However, empirical evidence shows the implied volatilities of traded options vary across strike prices and exhibit a smile or skewed shape across moneyness, the ratio between strike and underlying asset price; (2) The Black-Scholes model assumes that the stock price follows geometric Brownian motion, thus the risk-neutral probability density function (PDF) of the stock price is lognormal. However, researchers observe excess kurtosis and negative skewness of unconditional returns of the underlying security, which is inconsistent with the lognormality assumption. The first abnormality is related to the second one, since statistics such as skewness, and kurtosis can be derived if we know the entire risk-neutral PDF. Therefore, the central empirical issue in option pricing is what distributional hypothesis is consistent with underlying asset prices and traded option prices.

In this paper we are interested in estimating the implied risk-neutral moment generating function (MGF) of the underlying asset from option prices. Little attention has been paid to this area in the literature, in contrast to the interest and development in extracting the risk-neutral PDF from option prices (see Section 2 below and Taylor (2005) for an excellent review). We try to back out the risk-neutral MGF using the wavelet method and price options based on the pricing formula derived by Ma (2006b). Details of the MGF, the wavelet

method, and the option pricing formula of Ma (2006b) will be provided in the following sections. The contributions of this paper are as follows.

1. Although there is a one-to-one relationship between the MGF and the PDF, the MGF is more tractable in some cases. For instance, when there are random jumps in the price process, the PDF will not have a closed form while for the MGF we may expect an analytical expression (Ma and Vetzal, 1995).
2. The implied risk-neutral MGF obtained from our model is continuous while the implied risk-neutral PDF obtained from methods such as the smooth spline method is discrete.
3. With the estimated risk-neutral MGF, out-of-sample options with different time-to-maturity, different strike prices and even different underlying security prices can be calculated easily. This is contrary to existing risk-neutral distribution (RND) estimation methods that can only be used to infer the distribution for a specific time. Therefore, the estimated PDF can only be used to price out-of-sample options with a fixed expiry date. From a practical point of view, estimating the risk-neutral MGF is more appealing.
4. It is well known that option prices contain rich information on the implied volatility, the preference parameter, the jump process and the higher moments of the distribution. Based on the model of Ma (1992) and Ma (2006b), we are able to obtain the statistical moments of the underlying asset distribution direct from the risk-neutral MGF. The preference parameter of the utility function can also be revealed (Ma, 2006a).
5. There is no need for restrictions on the stochastic process of the under-

lying asset or prior assumptions on the implied risk-neutral MGF. This ensures flexibility of the method. Furthermore, wavelet functions which can represent any square integrable functions, are used to approximate the MGF. This is very flexible compared with the restrictive assumptions underlying a number of methods including the polynomial and the cubic spline method.

6. Contrary to the kernel estimation method of Ait-Sahalia and Lo (1998) or the neural network method of Hutchinson et al. (1994) and Garcia and Gençay (2000), the wavelet method does not require a large collection of data for a reasonable level of accuracy. We need only a small sample of options to estimate the MGF. For example, we can use only nine options with different strike prices on the same underlying asset with certain time-to-maturity to obtain a reasonably accurate risk-neutral MGF. We note that the kernel estimation and neural network methods require several thousand data points to obtain a reasonable level of accuracy.
7. Our technique avoids ill-posed inverse problems. According to Breeden and Litzenberger (1978), the risk-neutral PDF $g(X)$ can be obtained by differentiating the option pricing formula twice with respect to the strike prices X . For example, suppose there are three European call option prices c_1 , c_2 , and c_3 with time to maturity τ and strike prices of $K - \delta$, K , and $K + \delta$, respectively. Let the riskfree interest rate be r . The estimate of $g(X)$ at the point $X = K$ is given by $g(K) = e^{r\tau} \frac{c_1 + c_3 - 2c_2}{\delta^2}$ provided that δ is small enough. However, a major problem associated with this method is that the second derivative of the estimated call pricing function may not be a good estimator for the second derivative of the true function. This is because the option prices used for estimation are subject

to perturbations and small errors of option prices will be magnified when the denominator δ^2 is infinitely small. With the model in Ma (2006b), we avoid the problem by transforming it into a least square evaluation and estimating the parameters of a linear series which make up the risk-neutral PDF.

One may ask why the wavelet method is chosen instead of Fourier analysis. One of the reasons is that “in some cases (e.g. fingerprints) wavelet analysis is much better than Fourier analysis in the sense that fewer terms suffice to approximate certain functions” (Bachman et al., 2002, page 411). In addition, the Fourier series are a linear combination of a series of sine and cosine functions, which are defined over the entire real axis. Due to the periodic properties of the component sine and cosine functions, Fourier analysis is appealing in representing periodic functions. However, for non-periodic functions such as financial time series, the Fourier methodology is less suitable since there is little repetition within the sampled region. Therefore, wavelets, which are not restricted to a fixed shape or position, are more effective in dealing with non-periodic functions or non-stationary data series such as financial time series.

Wavelets as a mathematical analysis tool have been broadly applied in engineering including data-compression, de-noising, edge-detection, earthquake prediction, and so on. However, the use of wavelets in finance and economics is only a recent phenomenon. Despite this, wavelets are a very useful tool in financial and economics analysis as the examples in the following sections demonstrate.

The rest of the paper is organized as follows. The next section provides an overview of existing methods to reveal the risk-neutral PDF. We introduce wavelets in Section 3. Section 4 offers theoretical primitives on the risk-neutral MGF and wavelets. Section 5 describes the model and methodology of the

wavelet option pricing. Simulation and experimental results are given in section 6. Finally, Section 7 concludes.

2 Revealing the risk-neutral PDF

The existing methods for revealing the risk-neutral PDF can be grouped mainly into two categories: parametric and nonparametric ones. We start first with parametric methods.

1. The first parametric approach is to fit the option pricing function or the implied volatility smile parametrically. The risk-neutral PDF is then derived following the result of Breeden and Litzenberger (1978). To implement this approach, one needs strike prices varying from zero to infinity continuously. However, option contracts are only traded at discrete strike prices over a limited range. Therefore, most efforts have been focused on interpolating option prices within the range of strike prices and extrapolating outside the range to estimate the entire distribution. See for instance, Shimko (1993), Malz (1997), Bates (1991), and Bliss and Panigirtzoglou (2002).
2. An alternative approach is to specify a parametric stochastic process for the underlying asset price and the parameters can be recovered using observed market option prices. The risk-neutral PDF can then be inferred from the stochastic process. For instance, the classic Black and Scholes model assumes that the stock prices follow the Geometric Brownian motion with a constant risk-free rate and constant volatility implying a log-normal distribution for the stock prices. Other examples include Duffie et al. (2000) and Bates (2000) for the jump-diffusion process and stochastic

volatility model.

3. The third approach assumes a parametric form for the risk-neutral PDF of the underlying asset directly, and the parameters of the risk-neutral PDF can be estimated by minimizing the distance between the observed option prices and the fitted prices based on the model. For example, Melick and Thomas (1997) assume a mixture of three lognormal distributions for the implied PDF, and the estimation is carried out using the bounds on the American option prices. The method is claimed to be flexible, general and direct.

Another strand of the literature utilizes non-parametric methods. Non-parametric methods are considered more flexible since there is no prior restriction on the stochastic process of the underlying asset prices, the option pricing function, or the return distributions. For example, Rubinstein (1994) proposes establishing a prior distribution as a guess of the risk-neutral probabilities. The implied probability is inferred by minimizing the distance between the implied distribution and the prior distribution. The approach is non-parametric in that any probability distribution is a possible solution. This method requires a large amount of options so that the implied risk-neutral PDF is not dependent on the prior guess distribution. Ait-Sahalia and Lo (1998) provide an option pricing formula nonparametrically with kernel regression and the corresponding risk-neutral PDF is obtained numerically by differentiating option prices twice with respect to the strike price. Neural networks have also been used in nonparametric option pricing in Hutchinson et al. (1994), Garcia and Gençay (2000), and Gençay and Gibson (2007).

In summary, parametric methods need to make certain assumptions on the relations between variables or statistical properties of the asset return distribution.

This inevitably makes parametric methods more rigid. Nonparametric methods are often more flexible. However, nonparametric methods are usually data intensive, requiring a large sample of traded option prices in order to achieve a reasonable level of estimation accuracy. As will be discussed in the following sections, the nonparametric wavelet option pricing model that we propose is very flexible but needs only a small sample of option prices.

3 Wavelets: a brief overview

In this paper we propose an alternative to the methods reviewed above. We estimate the implied risk-neutral MGF instead of the risk-neutral PDF and we utilize wavelets to approximate the implied MGF function.

Although wavelets have not been widely applied in the financial and economic area, there is a growing literature in this regard. Ramsey (1999) provides an extensive review on applying the wavelet analysis to financial and economic data. The following are a few examples.

1. The first area of application is multi-resolution analysis or time scale analysis (also known as time-scale decomposition), which is powerful in revealing the relationship between economic variables and improving forecasts. An early key article is Davidson et al. (1998), in which the authors apply multi-resolution analysis on US commodity price behavior and obtain information on both the time location and the time scale of price movements. In the paper, the authors propose that wavelet analysis may help forecast price movements. This point was proven in Murtagh et al. (2004), whereby the authors examine several wavelet applications in time series prediction. After studying wavelet-based multiresolution autore-

gression models and single resolution approaches, the authors conclude that wavelet-based multiresolution approaches outperform the traditional single resolution approach in forecasting. Ramsey and Lampart (1998a) use wavelets to analyze the relationship between the expenditure and income at six different time scales and find that the relationship varies across time scales. Furthermore, the authors confirm that (page 23) “the time-scale decomposition is very important for analyzing economic relationships and that a number of anomalies previously noted in the literature are explained by these means”. Ramsey and Lampart (1998b) conduct a similar analysis on the relationship between money and domestic product. In a subsequent paper, Connor and Rossiter (2005) carry out a wavelet-based scale approach to analyze commodity prices motivated by the fact that “the dynamics of commodity markets have always been influenced by the interactions of traders with different time horizons, who react to the arrival of new information in a heterogeneous manner”. Mitra (2006) also exploits wavelets to perform multi-resolution analysis on the econometric relationship between money, output and price in the Indian macro economy. More examples can be found in Gençay et al. (2002), Capobianco (2002), and Yousefi et al. (2005).

2. The second area of application is de-noising financial time series data to reveal market trends. Gao and Ren (2005) use wavelets to analyze the highly erratic Shanghai Stock Market Index and consider it effective in suppressing the noise in the market index. Similarly, Antoniou and Vorlow (2005) apply wavelets to de-noise the FTSE-100 stock returns time series and find evidence of “non-periodic cyclical dynamics”. More examples can be found in Gençay et al. (2002).

3. The third application concerns the estimation of growth models and function approximations. For example, Esteban-Bravo and Vidal-Sanz (2007) propose a wavelet-based methodology for solving boundary value problems in growth models. Their results show that the wavelet method provides a very good approximation. Wavelets are also well known for their remarkable ability in approximating any function that belongs to $L^2(\mathbb{R})$. In Park et al. (2005), a wavelet-based Bayesian method is exploited in function estimation. The authors conclude that “the wavelet procedure appears to do a very good job at estimating both the function and the other parameters of the model, for all directions and noise levels considered in the study”. A comparison with other existing methods suggests that the wavelet-based Bayesian method outperforms the splines-based Bayesian approach in Antoniadis et al. (2004).

While wavelets have many useful properties, in this paper we are interested in one of the most basic features of the wavelets, i.e. function approximation. Each wavelet function has its characteristic and is suitable to approximate different types of functions. In our research, we consider a number of wavelets, including the Haar wavelet (Figure 1), the Franklin hat function (Figure 2), and the Shannon wavelet (Figure 3). We find that the Franklin hat function requires the smallest number of terms to approximate the MGF function at the same level of goodness of fit. Therefore, we use the Franklin hat function to derive the risk-neutral MGF from option prices. With the estimated MGF, we will further carry out out-of-sample tests to demonstrate the capability of the wavelet method in approximating functions of unknown forms.

4 Theoretical primitives

4.1 MGF

The moment generating function of a continuous random variable x is defined as the bilateral Laplace transform of the probability density function of the variable,

$$M(s) = \int_{-\infty}^{\infty} p(x)e^{-xs}dx, \quad (1)$$

where $p(x)$ is the PDF of x and s is a complex value in the complex plane. The PDF is uniquely determined by the inverse Laplace transform of the MGF. In the risk-neutral world, the same relationship also holds. The definitions of the Laplace transform and the inverse Laplace transform are provided in the Appendix.

4.2 The wavelet method

As the name suggests, a wavelet is a wave-like function. It is localized over a short interval and is zero except on that short interval. The wavelet function oscillates around its average value (zero) over a short distance, and damps out very fast outside the interval. A wavelet function must satisfy three criteria. (1) A wavelet must be square integrable; (2) the Fourier transform $\hat{\psi}(f)$ of the wavelet function $\psi(t)$ must satisfy the condition (Addison, 2002, page 9): $\int_0^{\infty} \frac{|\hat{\psi}(f)|^2}{f} df < \infty$; and (3) the wavelets integrate to zero, which ensures their oscillatory shape.

Unlike Fourier series, which have only sine and cosine basis functions, there is a large family of different wavelet basis functions. Consider the simple case of

a Haar wavelet. The Haar mother wavelet is defined as follows:

$$\psi(t) = \begin{cases} 1 & \text{if } 0 \leq t < \frac{1}{2} \\ -1 & \text{if } \frac{1}{2} \leq t < 1 \\ 0 & \text{otherwise} \end{cases} \quad (2)$$

The function looks like a square wave. It has non-zero values over the short interval $[0,1]$ and it disappears outside this range. Two types of manipulations can be performed on the mother wavelet to change its shape and position to generate other wavelets. The first type of manipulation is dilation (scaling), which squeezes or stretches a wavelet. The second type is translation, which shifts the wavelets horizontally. In Figure 1, the wavelet on the top is the mother Haar wavelet. In the middle row, the mother wavelet is squeezed to half and stretched to twice the original width, respectively. In the bottom row, the wavelet is shifted to the right and left, respectively. These wavelets are called generations of the mother wavelet $\psi(t)$. For any arbitrary wavelet function $\psi(\cdot) \in L^2(\mathbb{R})$, the generations $\psi_{l,k}(\cdot)$ are given by $\psi_{l,k}(t) \equiv 2^{\frac{l}{2}} \psi(2^l t - k)$, $l, k = 0, \pm 1, \pm 2, \dots$. The parameter l determines the level of dilation or contraction and the parameter k governs the movement of the wavelet along the horizontal axis.

The wavelet functions $\psi_{l,k}(t)$ can be chosen to be orthogonal to each other and are normalized. Therefore, $\psi_{l,k}(t)$ form an orthonormal wavelet basis for $L^2(\mathbb{R})$. Having defined the wavelet basis, we can represent any square integrable function $x(t)$ by adding up wavelet basis functions $\psi_{l,k}(t)$ over all integers l and k :

$$x(t) = \sum_{l=-\infty}^{\infty} \sum_{k=-\infty}^{\infty} T_{l,k} \psi_{l,k}(t); \quad (3)$$

where $T_{l,k}$ are the wavelet coefficients and they can be obtained through multiplying the function $x(t)$ with the basis function $\psi_{l,k}(t)$ as follows,

$$T_{l,k} = \int_{-\infty}^{\infty} x(t) \psi_{l,k}(t) dt. \quad (4)$$

It may seem challenging to estimate the unknown function $x(t)$ since we need to add up an infinite number of functions to avoid information loss. Fortunately, the coefficients $T_{l,k}$ converge to zero quickly as the parameters l and k increase so that we can truncate the coefficients at some point. All we need from the wavelet series is a good MGF approximation that captures enough information to allow us to price and forecast option prices.

As we have mentioned before, the Franklin hat function will be used in this paper. Although Franklin mother wavelets are complicated, they can be deduced from a simple hat function defined as follows,

$$h(t) = \begin{cases} (1 - |t|) & \text{if } -1 \leq t < 1 \\ 0 & \text{otherwise} \end{cases}, \quad (5)$$

with its Laplace transform $m_h(s) = (\frac{e^{s/2} - e^{-s/2}}{s})^2$. The dilated and translated versions of the hat function are given by $h_{l,k}(t) \equiv 2^{\frac{l}{2}} h(2^l t - k)$, $l, k = 0, \pm 1, \pm 2, \dots$. As a result, we can approximate function $x(t)$ in terms of the Franklin hat function $h_{l,k}(t)$. Note that the Franklin hat function does not integrate to 0.

There exists a huge family of wavelet functions. Most of the wavelet literature seems to concur upon the position that there is no best wavelet for a particular application. Most of the wavelet functions are probably able to approximate the MGF under a trade-off between the level of accuracy and the length of computation time. Here, we follow the advice of Mallat (1999) who remarks: “The design of (the wavelet) must... be optimized to produce a maximum number of wavelet coefficients that are close to zero.” (Mallat, 1999, page 241). We want fast computation since we deal with option prices and the Franklin hat function delivers on this criterion. Furthermore, it has the properties of being symmetric, smooth, and piecewise continuous, and it closely emulates the PDF

function of asset returns. We find that amongst the three types of wavelets we have mentioned, the Franklin hat function provides for the best trade-off between accuracy and computation intensity. For a detailed review on wavelets, see Hubbard (1998) and Chui (1992). A more formal and thorough treatment on wavelet theory can be found in Daubechies (1992) and Bachman et al. (2002).

5 The model and methodology

5.1 The model

The model in this paper is based on Ma (2006b). The author derives a closed form formula for European call options in a particular parameterization of the economy, which generalizes a number of option pricing models in the existing literature. Assuming that the moment generating function for the logarithm of time T asset price $\ln S_T$ is well defined, the option pricing formula is as follows,

$$C_t(S_t, X, T) = X e^{-r(T-t)} \mathcal{L}^{-1} \{ \Phi_{T-t}(s) \} \left(\ln \frac{X}{S_t} \right); \quad (6)$$

where t is the current time; the operation symbol $\mathcal{L}^{-1}$ denotes the bilateral inverse Laplace transform operator (see Appendix for more details); C_t is the time- t equilibrium price of the European call option; S_t is the underlying asset price at time t ; X is the strike price; T is the maturity date; r is the continuously compounded risk free interest rate; and $\Phi_{T-t}(s) \equiv \frac{\Theta^{T-t}(s)}{s(s+1)}$, where s is a complex value with $Re(s) \in (x^*, -1)$ and $-x^*$ is the highest defined statistical moment for the underlying asset distribution, and $\Theta^{T-t}(s)$ is the risk-neutral MGF of the logarithmic return $\ln \frac{S_T}{S_t}$. When $T - t = 1$, $\Theta(s)$ is the risk-neutral MGF for the rate of return per unit of time.

This model can be derived as follows (Ma, 2006b). Let y denote $\ln S_T$, $G(e^y)$

denote option payoff $G(e^y) = (e^y - X)^+$, and $p(y)$ denote the risk-neutral PDF for y . We have,

$$\begin{aligned}
& e^{r(T-t)} C_t(S_t, X, T) \\
&= \int_{\mathcal{R}} p(y) G(e^y) dy \\
&= \int_{\mathcal{R}} \mathcal{L}^{-1}\{S_t^{-s} \Theta^{T-t}(s)\}(y) G(e^y) dy \\
&= \int_{\ln X}^{\infty} (e^y - X) \left[\frac{1}{2\pi i} \int_{\sigma-i\infty}^{\sigma+i\infty} \Theta^{T-t}(s) e^{(y-\ln S_t)s} ds \right] dy \\
&= \frac{1}{2\pi i} \int_{\sigma-i\infty}^{\sigma+i\infty} S_t^{-s} \Theta^{T-t}(s) \left[\int_{\ln X}^{\infty} e^{sy} (e^y - X) dy \right] ds \\
&= \frac{1}{2\pi i} \int_{\sigma-i\infty}^{\sigma+i\infty} S_t^{-s} \Theta^{T-t}(s) \left[\frac{X^{s+1}}{s(s+1)} \right] ds \\
&= \frac{X}{2\pi i} \int_{\sigma-i\infty}^{\sigma+i\infty} \Phi_{T-t}(s) \left(\frac{X}{S_t} \right)^s ds \\
&= X \mathcal{L}^{-1}\{\Phi_{T-t}(s)\} \left(\ln \frac{X}{S_t} \right), \sigma \in (x^*, -1)
\end{aligned} \tag{7}$$

The second equality follows from the fact that the MGF for $y = \ln S_T$ is given by $S_t^{-s} \Theta^{T-t}(s)$, where $\Theta^{T-t}(s)$ is the MGF for $\ln \frac{S_T}{S_t}$. The sixth equality follows from the definition $\Phi_{T-t}(s) \equiv \frac{\Theta^{T-t}(s)}{s(s+1)}$.

Note that because the time to maturity $T - t$ is the power of the risk-neutral MGF $\Theta(s)$, it can be modified arbitrarily to $T' - t'$ for a different expiry date or different time to maturity in the out-of-sample forecast. In doing so, we assume that stock returns are independent and identically distributed (iid) and that volatility is constant. These assumptions are valid descriptions of the market for short time intervals. This flexibility is a unique feature of using the MGF representation for the option pricing formula. It cannot be exploited when the risk-neutral PDF approach is employed.

Note also the fifth equality in (7) is valid only for s with $Re(s) < -1$. If the real part of s is greater than -1, the inverse Laplace transform does not converge to the option prices.

As a special case, the Black-Scholes formula with a constant dividend-equity ratio l can be obtained by substituting the risk-neutral MGF

$$\Theta(s) = e^{-(r-l-\frac{\sigma^2}{2})s + \frac{\sigma^2}{2}s^2} \quad (8)$$

for the rate of annual return into equation (6), where r represents the drift, and σ stands for the volatility of the underlying stock price.

5.2 Methodology

In this section, we explain how we use Monte Carlo simulations to approximate the risk-neutral MGF with wavelets. Option prices are first simulated according to the benchmark Black-Scholes model. We approximate the risk-neutral MGF for the simulated option prices using the wavelet method outlined in Ma (2006b). The inferred MGF is then substituted into the option pricing function (6) so that we can compare the fitted option prices with those from the Black-Scholes model. Note that our model does not assume that the Black-Scholes model correctly describes the market. Instead the Black-Scholes model is employed as a benchmark to demonstrate the effectiveness of the wavelet method. Therefore, even if the traded option prices are determined by some unknown model, we can still back out the MGF with wavelets as we do not make prior assumptions.

The experiments are conducted as follows:

1. For an underlying asset with price S_t at time t , N different strike prices $X = \{X_1, X_2, \dots, X_N\}$ with a common time-to-maturity T , the riskfree interest rate r , and volatility σ , generate corresponding option prices $C^{bs} = \{C_1^{bs}, C_2^{bs}, \dots, C_N^{bs}\}$ using the Black-Scholes formula.
2. Given a set of scale and shift parameters, estimate the risk-neutral MGF $\hat{\Theta}(s)$ of the annual logarithmic return of the underlying asset from the data

set $\{S_t, X, T, C^{bs}\}$ using wavelet analysis. We calculate option prices using equation (6) with the derived MGF $\hat{\Theta}(s)$. Let $C^w = \{C_1^w, C_2^w, \dots, C_N^w\}$ denote the fitted wavelet-based option prices. Compare C^w with C^{bs} to assess the in-sample goodness of fit.

3. To test the out-of-sample forecast ability, select another data set $\{S'_t, X', T', C^{bs'}\}$. Calculate the wavelet based option prices $C^{w'}$ with the derived risk-neutral MGF $\hat{\Theta}(s)$ and compare $C^{w'}$ with $C^{bs'}$ to assess the out-of-sample forecast deviation.

Among the three steps above, Step 2 is the key so we discuss this step in detail below. The algorithm used to find the wavelet coefficients that fit the data C^w to C^{bs} is suggested in Ma (2006a). The following explains how we use wavelets to approximate the risk-neutral MGF.

As we have stated before, the mother wavelet function $\psi(\cdot) \in L^2(\mathbb{R})$ can be chosen such that its generation

$$\psi_{l,k}(x) \equiv 2^{\frac{l}{2}} \psi(2^l x - k), \quad l, k = 0, \pm 1, \pm 2, \dots \quad (9)$$

forms an orthonormal basis for $L^2(\mathbb{R})$. Let $m_\psi(s)$ and $m_{l,k}(s)$ denote the Laplace transform of $\psi(x)$ and $\psi_{l,k}(x)$ respectively, where $l, k = 0, \pm 1, \pm 2, \dots$. We have $m_{l,k}(s) = 2^{-\frac{l}{2}} e^{-\frac{ks}{2^l}} m_\psi(\frac{s}{2^l})$, $l, k = 0, \pm 1, \pm 2, \dots$

Assuming that the PDF $p(x)$ of a random variable x belongs to $L^2(\mathbb{R})$, we can expand $p(x)$ in terms of the orthonormal wavelet basis:

$$p(x) = \sum_l \sum_k a_{lk} \psi_{l,k}(x). \quad (10)$$

Perform a Laplace transform on both sides of equation (10) and we obtain $\Theta(s) = \sum_l \sum_k a_{lk} m_{l,k}(s)$, $Re(s) \in (x^*, 0]$; where $\Theta(s)$ is the risk-neutral MGF

of the random variable x and it is equal to the Laplace transform of the risk-neutral PDF $p(x)$.

To estimate the risk-neutral MGF with simulated data set $\{S_t, X_i, T, C_i^{bs}\}, i = 1, 2, \dots, N$, we may follow the procedures below,

1. For positive integers L and K , truncate the coefficients by setting $a_{lk} = 0$ for all $|l| > L$ and $|k| > K$. Set $\theta_{L,K} \equiv \{a_{lk}\}_{l \leq L, |k| \leq K}$.
2. Given the data set, we estimate the unknown coefficients $\theta_{L,K}$ by minimizing the average of squared errors between the true option prices C_i^{bs} and estimated prices C_i^w obtained by substituting $\hat{\Theta}(s|\theta_{L,K})$ into formula (6).
3. Go to step 1 with $L \rightarrow L + 1$ and $K \rightarrow K + 1$ until $\sum_i (C_i^{bs} - C_i^w)^2 < \varepsilon$, for any arbitrary $\varepsilon > 0$.

However, there are several issues in the above procedures. First, with increasing scale and shift parameters L and K , the computation time for iteration increases dramatically as the number of iterations increases geometrically. Second, the increasing number of parameters requires more data for the optimization. So during the iteration process it always runs out of data before reaching the optimum. Therefore, we fix the scale parameter l say at $l = L = 3$, and let shift parameters k change from $-K$ to K . We moderate Step 3 above accordingly into,

- Go to step 1 with $K \rightarrow K + 1$ until $\sum_i (C_i^{bs} - C_i^w)^2 < \varepsilon$ for any arbitrary $\varepsilon > 0$.
- If the fitting result improves very little with the increase of $K \rightarrow K + 1$ so that the optimization process does not terminate within a reasonable time

duration, increase $L \rightarrow L+1$ and repeat the above steps until a satisfactory result is obtained. This optimization process yields an estimate of the risk-neutral MGF

$$\hat{\Theta}(s) = \sum_{|l|=L} \sum_{|k| \leq K} \hat{a}_{lk} m_{lk}(s). \quad (11)$$

Now we face another issue of how to quickly and effectively search the scale and shift parameters l and k for the wavelet series. According to the relationship between the PDF and the MGF, the estimated coefficients $\hat{a}_{lk}$ and corresponding wavelet function can be used to form the risk-neutral PDF of the annual logarithmic return of the underlying asset. Therefore, we may be able to select the appropriate initial scale and shift parameters according to the interval of annual logarithmic returns $x = \ln(\frac{S_1}{S_0})$, which lies typically over the interval $[-0.7, 0.7]$. Wavelets have non-zero values for a small region. For the Franklin hat function it has a closed and bounded interval between $[-1, 1]$, and is zero otherwise. Hence, according to equation (9), for $\psi_{l,k}(x)$ to be effective in composing the probability density function, we need to calibrate the scale and shift parameters to ensure that $-1 \leq 2^l x - k \leq 1$. Assume that we have chosen a suitable scale parameter l , the shift parameters k should lie in the interval $[2^l x_{\min} - 1, 2^l x_{\max} + 1]$. We let $x_{\min} = -0.7$ and $x_{\max} = 0.7$ and these give

$$k \in [-0.7 \times 2^l - 1, 0.7 \times 2^l + 1]. \quad (12)$$

Note that the scale parameter determines the resolution of the estimated risk-neutral MGF. The larger the scale parameter, the finer the estimated risk-neutral MGF provided that the shift parameters are appropriately chosen. According to equation (12), on the one hand we need more shift parameters k to perform the approximation but this will be time consuming. On the other hand, we may obtain a feasible solution with least square estimation within several

minutes with a small scale parameter l . However, this is obtained at the cost of approximation accuracy.

6 Simulations and experimental results

We perform constrained least square estimation in this section. There are several restrictions on the call pricing function

$$C(S_t, X, \tau, r) = e^{-r\tau} \int_X^\infty (S_T - X) p(S_T | S_t, \tau, r) dS_T. \quad (13)$$

First, the probability density must be non-negative. Second, the integral of the probabilities over the possible terminal asset price should be one. Third, the call option pricing function should be monotonically decreasing with respect to the strike price, implying that the first derivative of the pricing function with respect to the strike price should be negative. And fourth, the call pricing function should be convex with respect to the strike prices, implying that the second derivative of the pricing function with respect to the strike price should be positive.

The optimization process with four restrictions requires long computing time. In our case, the latter two restrictions can be inferred from the first two as shown below. Differentiating the above call price function with respect to strike price

$$\frac{\partial C}{\partial X} = -e^{-r\tau} \int_X^\infty p(S_T | S_t, \tau, r) dS_T. \quad (14)$$

As

$$p(S_T | S_t, \tau, r) \geq 0 \quad (15)$$

and

$$\int_0^\infty p(S_T | S_t, \tau, r) dS_T = 1 \quad (16)$$

we have $0 \leq \int_X^\infty p(S_T|S_t, \tau, r) dS_T \leq 1$ and this leads to the third constraint

$$-e^{-r\tau} \leq \frac{\partial C}{\partial X} \leq 0. \quad (17)$$

Twice differentiate the call price function

$$\frac{\partial^2 C}{\partial X^2} = e^{-r\tau} p(X|S_t, \tau, r) \geq 0. \quad (18)$$

This is non-negative since both $e^{-r\tau}$ and $p(S_T|S_t, \tau, r)$ are non-negative and this leads to the fourth restriction.

To summarize, the first two constraints in equations (15) and (16) ensure that the PDF is non-negative and integrates to 1, respectively. They are also sufficient conditions to ensure monotonicity in equation (17) and convexity in equation (18) of the call price function. Therefore, we need to impose only the first two constraints on the wavelet estimator. We use a constrained optimizer in Matlab to find the minimum mean squared errors between the wavelet option prices based on equation (6) and simulated option prices based on the Black-Scholes model. We use initial value of 1 for all the coefficients $\hat{a}_{lk}$. The optimization converges rapidly. When we repeat the optimization using the estimated coefficients as initial values, it always converges to very similar values that typically differ from the third decimal point.

Below we carry out three experiments by generating historic option prices with time-to-maturity of one year, one month, and six months, respectively.

6.1 Time-to-maturity: one year

Let $\{X_1, X_2, \dots, X_N\}$ denote a strike price sample of size N . We assume that we observe nine call options written on a non-dividend-paying stock at $S = 100$, and the strike prices are evenly spaced between 80 and 120 with increment

of 5. Other variables include: time to maturity $T = 1$, risk-free interest rate $r = 0.0559$, and volatility $\sigma = 0.2$. Let C^{bs} and C^w denote respectively the true option prices based on the Black-Scholes model and the wavelet estimated option prices. Given the information above, we perform wavelet analysis to estimate the risk-neutral MGF. We find that we are able to estimate the MGF with only nine Franklin hat functions with scale parameter $l = 3$ and shift parameters varying from $k = -4 : 4$.

The pricing errors are reported in Table 1 Panel A. The mean squared errors between C^{bs} and C^w is 1.6046×10^{-5} and mean absolute pricing error is 0.0037. As a matter of fact, pricing errors reported in Panel A are all very small and close to zero. Figure 4 plots the true and fitted option prices. It is evident that the differences between the true and fitted option prices are so small that they cannot be distinguished easily. The wavelet coefficients are $\{\hat{a}_{lk}\} = [0.0482 \ 0.2034 \ 0.0554 \ 0.5405 \ 0.7500 \ 0.6005 \ 0.4971 \ 0.0220 \ 0.1112]$.

The Black-Scholes MGF $\Theta(s)$ and the approximated MGF $\hat{\Theta}(s)$ for the annual logarithmic return $\ln(\frac{S_1}{S_0})$ are plotted in Figure 5. This figure is produced by calculating the value of the functions for complex values s ranging from $-2 - 20i$ to $-2 + 20i$ with imaginary unit i as the increment. As we discussed in Section 5.1, the value of the real part of s does not matter as long as it is less than -1 to allow the inverse Laplace transform to converge. We plot the true and approximated MGF with the real part on the X-axis and the imaginary part on the Y-axis. Figure 5 shows that the shape of the estimated risk-neutral MGF is close to the true one.

The Black-Scholes model assumes that the stock prices follow a lognormal distribution. Therefore, we have the following distribution for the annual loga-

rithmic return $x = \ln(\frac{S_t}{S_0})$: $P(x) = \frac{1}{\sigma\sqrt{2\pi}}e^{-(x-(\mu-\frac{\sigma^2}{2}))^2/2\sigma^2}$. In Figure 6, we plot the lognormal distribution with $\mu = 0.0559$, and $\sigma = 0.2$ in the starred curve and the fitted risk-neutral PDF, obtained from coefficients $\{\hat{a}_{lk}\}$: $\hat{p}(x) = \sum_l \sum_k \hat{a}_{lk} \psi_{l,k}(x)$ in the dotted curve. Both densities integrate to one. To examine the goodness-of-fit of the estimated PDF, we carry out the Kolmogorov-Smirnov test between the true and fitted densities. The test statistic is 0.090 with p-value of 0.739 indicating that the difference between the two densities is not statistically significant.

The first derivatives of the approximated call pricing function with respect to the strike price are all negative and lie within the area $(-0.9, -0.2)$. Since $-e^{-r\tau} = -0.9456$ for $r = 0.0559$ and $\tau = 1$, the region that the first derivative falls in is in line with the constraint (equation (17)) above. The second derivative of the approximated call pricing function with respect to the strike price is positive, indicating that the call pricing function is convex.

Having obtained the estimated risk-neutral MGF, we are interested in pricing out-of-sample options. This is done for three groups of options. In the first group, options have different time-to-maturity from the in-sample ones. This takes advantage of the MGF representation of equation (6) as discussed in Section 5.1. In the second group, options have different time-to-maturity and strike prices from the in-sample options. In the third group, options have different time-to-maturity, strike prices and underlying stock price from the in-sample ones.

1. First, we use the estimated risk-neutral MGF to forecast out-of-sample options with different time-to-maturity. We use time-to-maturity of one month, three months, six months and nine months, respectively. We report

two types of forecasting errors between the true prices and the fitted prices. The mean squared error is computed as $(C_i^{bs} - \hat{C}_i^w)^2/N$, and the mean absolute error is calculated as $|C_i^{bs} - \hat{C}_i^w|/N$. Panel B in Table 1 reports the mean, minimum, maximum, and standard deviation of out-of-sample pricing errors. In Figure 7, we plot group of options with time-to-maturity of 84 days (3 months) as they have the biggest pricing errors. From both Panel B and Figure 7, we conclude that the estimated risk-neutral MGF is effective in pricing out-of-sample options with different time-to-maturity, especially for options with time-to-maturity close to the in-sample ones.

2. We test the capability of the fitted risk-neutral MGF to forecast option prices with different strike prices and different time-to-maturity. We choose the strike prices to be evenly spaced between 80 to 120 with step of 1. The pricing errors are tabulated in Panel C of Table 1. We can see that the maximum squared error of 0.1584 occurs when the time-to-maturity is 84 days and the true price is 6.8440. We plot this set of options in Figure 8. Again we are safe in saying that the revealed risk-neutral MGF performs well in forecasting out-of-sample options with different time-to-maturity and different strike prices.
3. We want to further stretch the out-of-sample robustness test by investigating the pricing errors when options have different time-to-maturity, strike price, and underlying stock price. We let the underlying asset price S_0 to be 120 and the strike prices are evenly spaced from 100 to 140 with an increment of 5. The time-to-maturity is assumed to be one month, three months, six months, and nine months. Panel D in Table 1 reports the out-of-sample forecasting errors. The largest mean squared error is 0.0517 and the largest mean absolute pricing error is 0.1970. Figure 9 plots the

option prices when time-to-maturity is 84 days, as this group of options have the largest mean squared error. From Panel D and Figure 9, we can see that the MGF can price out-of-sample options with different time-to-maturity, different strike price, and different underlying stock price with considerable accuracy.

In summary, the risk-neutral moment generating function estimated from options with one year of life can accurately forecast options with different time-to-maturity, different strike prices, and different underlying stock price. To complete the experiment, we now conduct two more out-of-sample estimations when the in-sample options have time-to-maturity of one month and six months, respectively.

6.2 Time-to-maturity: one and six months

Following section 6.1, we carry out two more estimations by changing the time-to-maturity of the in-sample options. We set the time-to-maturity to one month and six months, respectively. Other parameters remain unchanged. The estimated coefficients are $\{\hat{a}_{lk}\}=[0.0691 \ 0.1317 \ 0.2498 \ 0.4743 \ 0.6504 \ 0.6284 \ 0.4106 \ 0.1744 \ 0.0398]$ for the one-month options and $\{\hat{a}_{lk}\}=[0.0054 \ -0.0002 \ 0.2557 \ 0.5757 \ 0.7206 \ 0.6669 \ 0.3938 \ 0.1625 \ 0.0480]$ for the six-month options. Summary statistic of the pricing errors are reported in Table 2 and 3. In Panel A Table 2 for options with 1-month to maturity, the in-sample mean squared error is 0.0344 and the mean absolute error is 0.1604. In Panel A Table 3 for options with 6-month to maturity, the in-sample mean squared error is 0.0001 and the mean absolute pricing error is 0.0077. Compared with the results reported in Table 1 where the in-sample options have time-to-maturity of one year, the pric-

ing errors are slightly larger and the estimated risk-neutral MGF deviates from the true one slightly more. We suggest three reasons.

1. During the optimization process, the gradient is calculated in each iteration. However, when the time-to-maturity is not equal to 1, the gradient is more difficult to calculate as it takes longer for the optimization to converge and more function evaluations are needed. Therefore, larger errors are likely because the optimization may terminate when the maximum number of function evaluations is reached before reaching the global maximum value. In this case, we need to repeat the optimization process with the obtained coefficients as initial values. We repeat this process until we get a satisfactory result.
2. Equation (6) specifies that we have to carry out an inverse Laplace transform to estimate option prices. To do so, we need to integrate $\frac{\Theta^{T-t}(s)}{s(s+1)}$ from negative infinity to positive infinity, which cannot be achieved in practice. Therefore, we can integrate only over a truncated interval, which may lead to estimation errors.
3. In Table 2 where in-sample options have a time-to-maturity of one month and the strike prices vary from 80 to 120, four options have prices between zero and one. Options with prices near zero contain less information about the risk-neutral MGF than other options. This may result in a larger estimation error for the risk-neutral MGF.

Despite this, the estimated risk-neutral MGF still performs well overall in the out-of-sample forecast. In particular, the estimated risk-neutral MGF forecasts better when the out-of-sample options have time-to-maturity closer to the in-sample one. When the time-to-maturity increases, the forecast errors increase

slightly.

Our methodology also allows us to price other contingent claims written on the same underlying asset, including complex and illiquid derivative claims. This is because the risk-neutral MGF and PDF capture the same information and we can obtain the risk-neutral PDF from the estimated MGF. With the entire asset distribution and payoff structure, we will be able to price exotic derivatives.

7 Conclusions

In this paper, we apply the wavelet methodology for estimating the risk-neutral MGF of the underlying asset from options prices based on Ma (2006b). The most important contribution of our paper is that the wavelet method offers a promising alternative for pricing out-of-sample options and other complex and illiquid derivatives on the same underlying asset.

Our experiment contains three steps. First, we simulate a set of options prices using the benchmark Black-Scholes model. Second, we use a series of Franklin hat functions with different scale and shift parameters to estimate the implied risk-neutral MGF. Third, we compare option prices obtained from the fitted MGF with simulated option prices to see whether the wavelet method is effective in revealing the risk-neutral MGF. We also apply the fitted risk-neutral MGF to price out-of-sample options with different times-to-maturity, different strike prices, and different underlying asset prices. Through comparison between the option prices obtained from the fitted risk-neutral MGF and the simulated option prices, we provide strong evidence of the superior ability of the wavelet method in accurately approximating the risk-neutral MGF, and in pricing and forecasting option prices.

Acknowledgment

We thank the three referees and the editor for their detailed and helpful advice that greatly improved the paper. Thanks are also due to Jerry Coakley for his comments on the previous version of the paper. Chenghu Ma acknowledges funding from the Economic and Social Research Council (ESRC), UK. Liya Shen acknowledges funding from the Overseas Research Students Awards Scheme (ORSAS), UK. This paper is part of her PhD thesis.

Appendix: Laplace transform

For a function $f(t)$ which is real valued and piecewise continuous on $[0, \infty)$, its Laplace transformation is a complex valued function given by

$$\mathcal{L}\{f(t)\}(s) = F(s) = \int_0^{\infty} f(t)e^{-st} dt; \quad (19)$$

where s is a complex value in the complex plane and $\mathcal{L}$ denotes the Laplace transform operator. The inverse Laplace transform, denoted by $\mathcal{L}^{-1}\{F(s)\}(t)$, is defined as:

$$\mathcal{L}^{-1}\{F(s)\}(t) = f(t) = \frac{1}{2\pi i} \int_{c-i\infty}^{c+i\infty} F(s)e^{st} ds; \quad (20)$$

where c is a specific real number so that all singularities of $F(s)$ are to the left of it.

To introduce the Laplace transform into the option pricing model, we need not only positive t , but also negative t . Therefore, we need a so-called bilateral Laplace transform and bilateral inverse Laplace transform. The bilateral Laplace transformation of $f(t)$, denoted by $\mathcal{L}\{f(\cdot)\}(s)$, is given by

$$\mathcal{L}\{f(t)\}(s) = F(s) = \int_{-\infty}^{\infty} f(t)e^{-st} dt; \quad (21)$$

where $f(t)$ is defined for $t \in \mathbb{R}$, and s is a complex value in the complex plane. Let $F(s)$ denote $\mathcal{L}\{f(x)\}(s)$ and $G(s)$ denote $\mathcal{L}\{g(x)\}(s)$, we have the properties of the Laplace transform summarized as following:

1. Linearity

$$\mathcal{L}\{af(x) + bg(x)\}(s) = aF(s) + bG(s); \quad (22)$$

$$\mathcal{L}^{-1}\{aF(s) + bG(s)\}(x) = af(x) + bg(x). \quad (23)$$

2. Frequency shifting

$$\mathcal{L}\{e^{-lx}f(x)\}(s) = F(s+l), \forall l \in \mathbb{R}; \quad (24)$$

$$\mathcal{L}^{-1}\{F(s+l)\}(x) = e^{-lt}f(x), \forall l \in R. \quad (25)$$

3. Time shifting

$$\mathcal{L}\{f(x-x_0)\}(s) = e^{-x_0s}F(s), \forall x_0 \in R; \quad (26)$$

$$\mathcal{L}^{-1}\{e^{-x_0s}F(s)\}(x) = f(x-x_0), \forall x_0 \in R. \quad (27)$$

4. Convolution

$$\mathcal{L}\{f(x) * g(x)\} = F(s)G(s); \quad (28)$$

$$\mathcal{L}^{-1}\{F(s)G(s)\}(x) = f(x) * g(x). \quad (29)$$

where ‘*’ indicates the convolution operator on f and g . This operator can be defined as (Bracewell (1999, page 25)),

$$f * g \equiv \int_{-\infty}^{\infty} f(\tau)g(t-\tau)d\tau = \int_{-\infty}^{\infty} g(\tau)f(t-\tau)d\tau. \quad (30)$$

References

- Addison, P. S., 2002. The Illustrated Wavelet Transform Handbook: Introductory Theory and Applications in Science, Engineering, Medicine and Finance. IOP Publishing, Bristol.
- Ait-Sahalia, Y., Lo, A., 1998. Nonparametric estimation of state-price densities implicit in financial asset prices. *Journal of Finance* 52, 499–547.
- Antoniadis, A., Gregoire, G., McKeague, I. W., 2004. Bayesian estimation in single-index models. *Statistica Sinica* 14, 1147–1164.
- Antoniou, A., Vorlow, C. E., 2005. Price clustering and discreteness: is there chaos behind the noise? *Physica A* 348, 389–403.
- Bachman, G., Narici, L., Beckenstein, E., 2002. Fourier and Wavelet Analysis. Springer.
- Bates, D., 1991. The crash of 87': Was it expected? The evidence from options markets. *Journal of Finance* 46, 1009–1044.
- Bates, D., 2000. Post-87' crash fears in the S&P 500 futures options. *Journal of Econometrics* 94, 181–238.
- Black, F., Scholes, M., 1973. The pricing of options and corporate liabilities. *Journal of Political Economy* 81, 637–659.
- Bliss, R. R., Panigirtzoglou, N., 2002. Testing the stability of implied probability density functions. *Journal of Banking and Finance* 26, 381–422.
- Bracewell, R., 1999. The Fourier Transform and Its Applications. McGraw-Hill, US.

- Breeden, D. T., Litzenberger, R. H., 1978. Prices of state-contingent claims implicit in option prices. *Journal of Business* 51, 621–651.
- Capobianco, E., 2002. Independent component analysis and resolution pursuit with wavelet and cosine packets. *Neurocomputing* 48, 779–806.
- Chui, C. K., 1992. *An Introduction to Wavelets. Vol. 1 (Wavelet Analysis and Its Applications)*. Academic Press, New York.
- Connor, J., Rossiter, R., 2005. Wavelet transforms and commodity prices. *Studies in Nonlinear Dynamics and Econometrics* 9 (1), Article 6.
- Daubechies, I., 1992. *Ten Lectures on Wavelets*. SIAM, Philadelphia.
- Davidson, R., Labys, W. C., Lesourd, J. B., 1998. Wavelet analysis of commodity price behavior. *Computational Economics* 11, 103–28.
- Duffie, D., Pan, J., Singleton, K., 2000. Transform analysis and asset pricing for affine jump-diffusions. *Econometrica* 68, 1343–1376.
- Esteban-Bravo, M., Vidal-Sanz, J. M., 2007. Computing continuous-time growth models with boundary conditions via wavelets. *Journal of Economic Dynamics and Control* 31, 3614–3643.
- Gao, L., Ren, H., 2005. Detect market baseline trend with wavelet analysis. *Wavelet Analysis and Active Media Technology* 1-3, pp. 1497–1502.
- Garcia, R., Gençay, R., 2000. Pricing and hedging derivative securities with neural networks and a homogeneity hint. *Journal of Econometrics* 94, 93–115.
- Gençay, R., Gibson, R., 2007. Model risk for European-style stock index options. *IEEE NN* 18, 193–202.

- Gençay, R., Selçuk, F., Whitcher, B., 2002. An Introduction to Wavelets and Other Filtering Methods in Finance and Economics. Academic press, New York.
- Hubbard, B. B., 1998. The World According to Wavelets, 2nd Edition. A K Peters, Natick.
- Hutchinson, J. M., Lo, A. W., Poggio, T., 1994. A nonparametric approach to the pricing and hedging of derivative securities via learning networks. *Journal of Finance* 49, 851–889.
- Ma, C., 1992. Two essays on equilibrium asset pricing and intertemporal recursive utility. Ph.D. thesis, Department of Economics, University of Toronto.
- Ma, C., 2006a. Advanced Asset Pricing Theory. Unpublished manuscript.
- Ma, C., 2006b. Intertemporal recursive utility and an equilibrium asset pricing model in the presence of levy jumps. *Journal of Mathematical Economics* 42, 131–160.
- Ma, C., Vetzal, K. R., 1995. Pricing options on the market portfolio with discontinuous returns under recursive utility. Mimeo, University of Waterloo.
- Mallat, S., 1999. A Wavelet Tour of Signal Processing. Academic Press, New York.
- Malz, A. M., 1997. Estimating the probability distribution of the future exchange rate from options prices. *Journal of Derivatives*, 18–36.
- Melick, W. R., Thomas, C. P., 1997. Recovering an asset’s implied pdf from option prices: an application to crude oil during the gulf crisis. *Journal of Financial and Quantitative Analysis* 32, 91–115.

- Mitra, S., 2006. A wavelet filtering based analysis of macroeconomic indicators: the indian evidence. *Applied Mathematics and Computation* 175, 1055–1079.
- Murtagh, F., Starck, J. L., Renaud, O., 2004. On neuro-wavelet modeling. *Decision Support Systems* 37, 475–484.
- Park, C. G., Vannucci, M., Hart, J. D., 2005. Bayesian methods for wavelet series in single-index models. *Journal of Computational and Graphical Statistics* 14, 770–794.
- Ramsey, J. B., 1999. The contribution of wavelets to the analysis of economic and financial data. *Philosophical Transaction of the Royal Society London A* 357, 2593–2606.
- Ramsey, J. B., Lampart, C., 1998a. Decomposition of economic relationship by time scale using wavelets: expenditure and income. *Studies in Nonlinear Dynamics & Econometrics* 3, 23–42.
- Ramsey, J. B., Lampart, C., 1998b. Decomposition of economic relationship by time scale using wavelets: money and income. *Macroeconomic Dynamics* 2, 49–71.
- Rubinstein, M., 1994. Implied binomial trees. *Journal of Finance* 49, 771–818.
- Shimko, D., 1993. Bounds of probability. *Risk* 6, 33–37.
- Taylor, S. J., 2005. *Asset Price Dynamics, Volatility, and Prediction*. Princeton University Press, Princeton.
- Yousefi, S., Weinreich, I., Reinartz, D., 2005. Wavelet-based prediction of oil prices. *Chaos Solitons & Fractals* 25, 265–275.

Figure 1. Dilation and translation to the Haar mother wavelet

Figure 2. Franklin hat function

Figure 3. Shannon wavelet

Figure 4. True and fitted options

Figure 5. True and fitted MGF

Figure 6. True and fitted PDF

Figure 7. Out-of-sample forecast (1)

Figure 8. Out-of-sample forecast (2)

Figure 9. Out-of-sample forecast (3)

Table 1. In-sample and out-of-sample pricing errors for options with one-year to maturity.

Panel A reports in-sample pricing errors. Panel B reports out-of-sample forecast for options with different time to maturity. Panel C reports out-of-sample forecast for options with different time to maturity and strike price. Panel D reports out-of-sample forecast for options with different time to maturity, strike price, and underlying asset price.

T	c_{BS}	mean	min	max	std deviation
Panel A. In-sample fit					
Squared errors					
365	12.2257	1.6046×10^{-5}	2.3132×10^{-6}	3.8171×10^{-5}	1.2081×10^{-5}
Absolute errors					
365	12.2257	0.0037	0.0015	0.0062	0.0015
Panel B. Out-of-sample forecast (T)					
Squared errors					
28	6.1368	0.0341	0.0017	0.1489	0.0518
84	7.2201	0.0449	0.0048	0.1584	0.0496
168	8.8083	0.0239	0.0012	0.0879	0.0277
252	10.3156	0.0054	4.5616×10^{-5}	0.0250	0.0081
Absolute errors					
28	6.1368	0.1425	0.0415	0.3859	0.1247
84	7.2201	0.1859	0.0692	0.3980	0.1079
168	8.8083	0.1322	0.0346	0.2965	0.0848
252	10.3156	0.0573	0.0068	0.1583	0.0491
Panel C. Out-of-sample forecast (T, K)					
Squared errors					
28	5.6946	0.0293	3.7812×10^{-6}	0.1489	0.0407
84	6.8440	0.0407	2.2375×10^{-5}	0.1584	0.0419
168	8.5008	0.0215	2.1009×10^{-5}	0.0879	0.0235
252	10.0524	0.0047	4.8679×10^{-7}	0.0250	0.0067
Absolute errors					
28	5.6946	0.1310	0.0019	0.3859	0.1114
84	6.8440	0.1739	0.0047	0.3980	0.1036
168	8.5008	0.1249	0.0046	0.2965	0.0780
252	10.0524	0.0543	0.0007	0.1583	0.0429
Panel D. Out-of-sample forecast (T, K, S)					
Squared errors					
28	6.3313	0.0355	1.9469×10^{-4}	0.1500	0.0475
84	7.7913	0.0517	0.0016	0.1402	0.0512
168	9.8596	0.0272	8.2177×10^{-6}	0.0846	0.0286
252	11.7720	0.0057	2.1060×10^{-5}	0.0242	0.0077
Absolute errors					
28	6.3313	0.1539	0.0140	0.3873	0.1156
84	7.7913	0.1970	0.0402	0.3744	0.1206
168	9.8596	0.1416	0.0029	0.2908	0.0899
252	11.7720	0.0612	0.0046	0.1556	0.0464

Table 2. In-sample and out-of-sample pricing errors for options with one-month to maturity

Panel A reports in-sample pricing errors. Panel B reports out-of-sample forecast for options with different time to maturity. Panel C reports out-of-sample forecast for options with different time to maturity and strike price. Panel D reports out-of-sample forecast for options with different time to maturity, strike price, and underlying asset price.

T	c_{BS}	mean	min	max	std deviation
Panel A. In-sample fit					
Squared errors					
28	6.1368	0.0344	0.0017	0.0860	0.0324
Absolute errors					
28	6.1368	0.1604	0.0408	0.2932	0.0987
Panel B. Out-of-sample forecast (T)					
Squared errors					
84	7.2201	0.0073	$1.1401 \cdot 10^{-4}$	0.0257	0.0089
168	8.8083	0.0081	$7.0984 \cdot 10^{-6}$	0.0238	0.0079
252	10.3156	0.0185	$6.4067 \cdot 10^{-4}$	0.1071	0.0347
365	12.2257	0.0573	$7.7899 \cdot 10^{-5}$	0.3103	0.1060
Absolute errors					
84	7.2201	0.0684	0.0107	0.1604	0.0545
168	8.8083	0.0769	0.0027	0.1542	0.0497
252	10.3156	0.1002	0.0253	0.3273	0.0978
365	12.2257	0.1512	0.0088	0.5570	0.1967
Panel C. Out-of-sample forecast (T, K)					
Squared errors					
84	6.8440	0.0073	$3.1225 \cdot 10^{-6}$	0.0261	0.0087
168	8.5008	0.0072	$3.2518 \cdot 10^{-7}$	0.0238	0.0064
252	10.0524	0.0141	$3.3319 \cdot 10^{-7}$	0.1071	0.0245
365	12.0025	0.0457	$7.3917 \cdot 10^{-6}$	0.3103	0.0819
Absolute errors					
84	6.8440	0.0694	0.0018	0.1616	0.0508
168	8.5008	0.0734	$5.7025 \cdot 10^{-4}$	0.1542	0.0426
252	10.0524	0.0892	$5.7722 \cdot 10^{-4}$	0.3273	0.0791
365	12.0025	0.1359	0.0027	0.5570	0.1671
Panel D. Out-of-sample forecast (T, K, S)					
Squared errors					
84	7.7913	0.0107	$1.8593 \cdot 10^{-4}$	0.0376	0.0135
168	9.8596	0.0095	$1.6785 \cdot 10^{-4}$	0.0230	0.0085
252	11.7720	0.0149	$8.6825 \cdot 10^{-5}$	0.0745	0.0232
365	14.1573	0.0511	$1.6016 \cdot 10^{-4}$	0.2743	0.0935
Absolute errors					
84	7.7913	0.0847	0.0136	0.1940	0.0631
168	9.8596	0.0849	0.0130	0.1515	0.0508
252	11.7720	0.0976	0.0093	0.2729	0.0780
365	14.1573	0.1471	0.0127	0.5238	0.1821

Table 3. In-sample and out-of-sample pricing errors for options with six-month to maturity.

Panel A reports in-sample pricing errors. Panel B reports out-of-sample forecast for options with different time to maturity. Panel C reports out-of-sample forecast for options with different time to maturity and strike price. Panel D reports out-of-sample forecast for options with different time to maturity, strike price, and underlying asset price.

T	c_{BS}	mean	min	max	std deviation
Panel A. In-sample fit					
Squared errors					
168	8.8083	0.0001	0.0000	0.0004	0.0001
Absolute errors					
168	8.8083	0.0077	0.0012	0.0207	0.0062
Panel B. Out-of-sample forecast (T)					
Squared errors					
28	6.1368	0.0212	0.0005	0.0600	0.0214
84	7.2201	0.0007	0.0000	0.0017	0.0007
252	10.3156	0.0030	0.0000	0.0221	0.0072
365	12.2257	0.0198	0.0000	0.1089	0.0364
Absolute errors					
28	6.1368	0.1271	0.0215	0.2449	0.0758
84	7.2201	0.0215	0.0032	0.0408	0.0159
252	10.3156	0.0335	0.0015	0.1486	0.0464
365	12.2257	0.0943	0.0033	0.3300	0.1108
Panel C. Out-of-sample forecast (T, K)					
Squared errors					
28	5.6946	0.0224	0.0003	0.0654	0.0192
84	6.8440	0.0007	0.0000	0.0017	0.0006
252	10.0524	0.0001	0.0000	0.0004	0.0001
365	12.0025	0.0155	0.0000	0.1089	0.0272
Absolute errors					
28	5.6946	0.1329	0.0164	0.2558	0.0696
84	6.8440	0.0220	0.0015	0.0418	0.0143
252	10.0524	0.0069	0.0000	0.0207	0.0060
365	12.0025	0.0845	0.0001	0.3300	0.0926
Panel D. Out-of-sample forecast (T, K, S)					
Squared errors					
28	6.3313	0.0341	0.0002	0.0864	0.0296
84	7.7913	0.0010	0.0000	0.0024	0.0009
252	11.7720	0.0015	0.0000	0.0099	0.0032
365	14.1573	0.0171	0.0000	0.0880	0.0295
Absolute errors					
28	6.3313	0.1581	0.0150	0.2939	0.1011
84	7.7913	0.0266	0.0038	0.0492	0.0169
252	11.7720	0.0266	0.0035	0.0993	0.0299
365	14.1573	0.0907	0.0030	0.2967	0.0998