ANZIAM J. 57(2016), 319–338
doi:10.1017/S1446181115000310

HISTORICAL BACKTESTING OF LOCAL VOLATILITY MODEL USING AUD/USD VANILLA OPTIONS

T. G. LING[1] and P. V. SHEVCHENKO[✉2]

(Received 1 November, 2014; accepted 29 May, 2015; first published online 17 February 2016)

Abstract

The *local volatility model* is a well-known extension of the Black–Scholes constant volatility model, whereby the volatility is dependent on both time and the underlying asset. This model can be calibrated to provide a perfect fit to a wide range of implied volatility surfaces. The model is easy to calibrate and still very popular in foreign exchange option trading. In this paper, we address a question of validation of the local volatility model. Different stochastic models for the underlying asset can be calibrated to provide a good fit to the current market data, which should be recalibrated every trading date. A good fit to the current market data does not imply that the model is appropriate, and historical backtesting should be performed for validation purposes. We study delta hedging errors under the local volatility model using historical data from 2005 to 2011 for the AUD/USD implied volatility. We performed backtests for a range of option maturities and strikes using sticky delta and theoretically correct delta hedging. The results show that delta hedging errors under the standard Black–Scholes model are no worse than those of the local volatility model. Moreover, for the case of in- and at-the-money options, the hedging error for the Black–Scholes model is significantly better.

2010 *Mathematics subject classification*: 91G20.

Keywords and phrases: local volatility model, backtesting, model validation, foreign exchange options, implied volatility.

1. Introduction

Under the well-known Black–Scholes pricing model [1], the asset price S is modelled with geometric Brownian motion,

$$dS_t = \mu S_t\,dt + \sigma S_t\,dW_t,$$

where μ is the drift, σ is the volatility and W_t is a standard Brownian motion. One of the key assumptions of this model is the *no-arbitrage condition*, which means that it

[1]University of Technology, Sydney, NSW, Australia; e-mail: Timothy.Ling-1@uts.edu.au.
[2]CSIRO Computational Informatics and University of Technology Sydney, NSW, Australia; e-mail: Pavel.Shevchenko@csiro.au.

is impossible to make a riskless profit. From this assumption it can be shown that the fair price of a derivative security with underlying asset S is equal to the mathematical expectation of the discounted payoff of the derivative. This expectation is computed with respect to a so-called *risk-neutral probability measure*. Furthermore, under this measure the dynamics of the asset price S is given by

$$dS_t = (r - q)S_t\,dt + \sigma S_t\,dB_t,$$

where B_t is a standard Brownian motion under the risk-neutral measure, r is the constant risk-free interest rate and q denotes the constant dividend yield. Given the price of an European option, strike, maturity and interest rates, the asset price volatility σ can be computed numerically from the Black–Scholes pricing formula. We say that σ is the volatility *implied* by the market price. If the Black–Scholes model was a perfect representation of the market, then the implied volatility would be equal for all market-traded options. This is definitely not the case in practice.

The implied volatility is heavily dependent on the strike price and maturity of the option. The *local volatility* model is an extension of the Black–Scholes framework which can account for this dependence, and it does so by making volatility a function of the current time and current spot price, that is, $\sigma(S_t, t)$ (see, for example, the articles by Derman and Kani [3] and Dupire [4]). Then the risk-neutral process is given by

$$dS_t = (r - q)S_t\,dt + \sigma(S_t, t)S_t\,dW_t.$$

In this paper, we provide results from historical backtesting of delta hedging errors under the local volatility and Black–Scholes frameworks, using Australian dollar or US dollar (AUD/USD) data from 2005 to 2011. To our knowledge, there has been no paper published discussing the results of empirical testing with real financial data. We note a related paper discussing local volatility in the context of foreign exchange (FX) markets using stochastic interest rates [2].

2. The Black–Scholes setup

Let $V(S, t)$ denote the *discounted* price of a contingent claim at time t with underlying asset price $S(t)$. Also, let r denote the risk-free interest rate, q the dividend yield and σ the asset price volatility. Within the Black–Scholes framework, V satisfies the fundamental partial differential equation (PDE)

$$\frac{\partial V}{\partial t} + \frac{1}{2}\sigma^2 S^2 \frac{\partial^2 V}{\partial S^2} + (r - q)S\frac{\partial V}{\partial S} - rV = 0$$

with known boundary condition $V(S(T), T)$ for the case of European options. This PDE can be obtained by applying a trading strategy, called *delta hedging*.

For ease of numerical implementation, we transform the above PDE with $X = \log S$. Routine calculations show that the transformed PDE is

$$\frac{\partial V}{\partial t} + \left(r - q - \frac{1}{2}\sigma^2\right)\frac{\partial V}{\partial X} + \frac{\sigma^2}{2}\frac{\partial^2 V}{\partial X^2} - rV = 0. \tag{2.1}$$

https://doi.org/10.1017/S1446181115000310 Published online by Cambridge University Press

2.1. The Black–Scholes pricing formula Assuming constant interest rates r_c, dividend yields q_c and volatility σ_c, the Black–Scholes formula for the price of an European call option with strike K and maturity T, at time $t = 0$, is given by

$$C(\sigma_c, K, T, r_c, q_c, S_0) = e^{-r_c T}(S_0 e^{(r_c - q_c)T}\Phi(d_1) - K\Phi(d_1 - \sigma_c\sqrt{T})), \tag{2.2}$$

where $\Phi(\cdot)$ denotes the cumulative distribution function for the standard normal distribution and

$$d_1 = \frac{\ln(S_0/K) + (r_c - q_c + \sigma_c^2/2)T}{\sigma_c\sqrt{T}}.$$

For the case where interest rates, dividend yields and volatility are time dependent, the Black–Scholes formula can be applied with the following substitutions [13]:

$$r_c = \frac{1}{T}\int_0^T r(u)\,du, \quad q_c = \frac{1}{T}\int_0^T q(u)\,du, \quad \sigma_c^2 = \frac{1}{T}\int_0^T \sigma^2(u)\,du, \tag{2.3}$$

where $r(t)$, $q(t)$ and $\sigma(t)$ are called the instantaneous interest rate, instantaneous dividend yield and instantaneous volatility, respectively. The interpretation is that in a small interval of time $[t, t + \Delta t]$, the amount of interest accrued (owed) is $r(t)\Delta t$. Note that instantaneous $r(t)$, $q(t)$ and $\sigma(t)$ are not observable in the market: instead, we observe the market interest rate yields and implied volatilities for different maturities corresponding to integral quantities in (2.3); see Sections 5.3 and 5.4 for more discussion on this point and for how to input the correct values from market data.

3. Local volatility

The local volatility model extends the Black–Scholes framework by making volatility a function of current asset price and time. In addition, we introduce time dependence for the interest rate and dividend yield. This leads to the following modification of equation (2.1):

$$\frac{\partial V}{\partial t} + \Big(r(t) - q(t) - \frac{1}{2}[\sigma(e^{X_t}, t)]^2\Big)\frac{\partial V}{\partial X} + \frac{1}{2}[\sigma(e^{X_t}, t)]^2\frac{\partial^2 V}{\partial X^2} - r(t)V = 0, \tag{3.1}$$

where the *local volatility*

$$\sigma(K, T) = \sqrt{\frac{2\theta T(\partial\theta/\partial T) + \theta^2 + 2[r(T) - q(T)]K\theta T(\partial\theta/\partial K)}{[1 + d_1 K\sqrt{T}(\partial\theta/\partial K)]^2 + K^2\theta T[(\partial^2\theta/\partial K^2) - d_1(\partial\theta/\partial K)^2\sqrt{T}]}}, \tag{3.2}$$

$$d_1 = \frac{\ln(S_0/K) + \int_0^T [r(t) - q(t)]\,dt + \theta^2 T/2}{\theta\sqrt{T}}.$$

Here, we define $\theta(K, T) = C^{-1}(V)$, where $V = C(\theta, \cdot)$ is given by the Black–Scholes formula (2.2), that is, θ is the market *implied volatility* for a vanilla option with strike K and maturity T. Equation (3.2) is called the *Dupire formula* [4] (see the technical report by Shevchenko [11, p. 49] for a proof of the formula in this particular form).

https://doi.org/10.1017/S1446181115000310 Published online by Cambridge University Press

TABLE 1. An example of AUD/USD market implied volatilities on 12 April 2005. The spot price for that day was $S_0 = 0.7735$.

Maturity/Δ	10ΔPut	25ΔPut	ATM	25ΔCall	10ΔCall
1 week	9.963%	9.088%	8.450%	8.213%	8.338%
1 month	10.913%	10.038%	9.400%	9.163%	9.288%
2 months	11.363%	10.488%	9.850%	9.613%	9.738%
3 months	11.713%	10.838%	10.200%	9.963%	10.138%
6 months	12.155%	11.280%	10.630%	10.430%	10.605%
1 year	12.400%	11.525%	10.850%	10.675%	10.850%
2 years	12.157%	11.350%	10.750%	10.650%	10.844%
3 years	12.013%	11.250%	10.700%	10.650%	10.888%
4 years	11.966%	11.225%	10.700%	10.675%	10.935%
5 years	11.819%	11.100%	10.600%	10.600%	10.881%

The functions $r(t)$ and $q(t)$ are the instantaneous rates; see Section 5.3 on how to determine these functions from market data. To compute local volatility, we require an implied volatility surface that can be interpolated from market data. There is no universal way to perform this interpolation. We now describe a simple method that yields good results for FX data.

3.1. Interpolating market implied volatility To compute the local volatility function (3.2), we need partial derivatives of the implied volatility surface $\theta(K, T)$. In practice, we only have a finite number of market data points, typically five values for a given maturity and about 10 maturities (see Table 1). We need some interpolating procedure for θ. This is an ill-posed problem, and there are a number of ways to interpolate these data points (see, for instance, the articles by Feil et al. [5] and White [12]). We use natural cubic splines to interpolate across strikes and maturities. This of course is not the only way to perform such an interpolation but it results in a very good fit, as seen in Table 2.

3.2. Our method to compute local volatility Suppose that we have market data for N different maturities, and that options are available for each maturity M. Let $K_j^{(i)}$ and $\theta_j^{(i)}$ denote the strike and implied volatility of the jth vanilla option with maturity T_i, respectively.

(1) *Interpolation across strikes:* For each market maturity T_i ($i \in \{1, \dots, N\}$), fit a natural cubic spline y_i through

$$(K_1^{(i)}, \theta_1^{(i)}), (K_2^{(i)}, \theta_2^{(i)}), \dots, (K_M^{(i)}, \theta_M^{(i)}).$$

Note that $y_i'(K) = (\partial\theta/\partial K)|_{(K,T_i)}$ and $y_i''(K) = (\partial^2\theta/\partial K^2)|_{(K,T_i)}$.

https://doi.org/10.1017/S1446181115000310 Published online by Cambridge University Press

TABLE 2. Average of absolute calibration errors from historical data (%). The value 0.005% appears repeatedly, but this is due to rounding of the results to three decimal places.

Maturity/Δ	10ΔPut	25ΔPut	ATM	25ΔCall	10ΔCall
1 week	0.005	0.005	0.005	0.005	0.005
1 month	0.005	0.005	0.005	0.005	0.005
2 months	0.005	0.005	0.005	0.005	0.005
3 months	0.005	0.005	0.005	0.005	0.005
6 months	0.005	0.005	0.005	0.005	0.005
1 year	0.005	0.005	0.005	0.005	0.005
2 years	0.005	0.005	0.005	0.005	0.005
3 years	0.005	0.005	0.005	0.005	0.005
4 years	0.005	0.005	0.005	0.005	0.005
5 years	0.005	0.005	0.005	0.005	0.006

(2) *Interpolation across maturities:* To find $\partial\theta/\partial K$ at any given (K, T), fit another natural cubic spline z through

$$(T_1, y_1'(K)), (T_2, y_2'(K)), \ldots (T_N, y_N'(K));$$

then $\partial\theta/\partial K = z(T)$.

(3) Similarly, to find $\partial^2\theta/\partial K^2$ at any given (K, T), fit another natural cubic spline w through

$$(T_1, y_1''(K)), (T_2, y_2''(K)), \ldots (T_N, y_N''(K));$$

then $\partial^2\theta/\partial K^2 = w(T)$.

(4) To find θ and $\partial\theta/\partial T$ at (K, T), fit a natural cubic spline u through

$$(T_1, y_1(K)), (T_2, y_2(K)), \ldots (T_N, y_N(K));$$

then $\theta(K, T) = u(T)$ and $\partial\theta/\partial T = u'(T)$.

(5) Substitute the above calculated θ, $\partial\theta/\partial T$, $\partial\theta/\partial K$ and $\partial^2\theta/\partial K^2$ into (3.2) and compute $[\sigma(K, T)]^2$. If $[\sigma(K, T)]^2 < 0$; then we overwrite $\sigma(K, T) = 0$.

Note that this method can obtain a value for local volatility for any (K, T) pair beyond the market range (for T smaller than the first market maturity, larger than the last market maturity etc.), by linear extrapolation of the natural cubic splines. For example, if we have a natural cubic spline $y(x)$ fitted to data points $x_1, x_2, \ldots, x_n$, then our function including extrapolation is given by

$$y^*(x) = \begin{cases} y'(x_1)(x - x_1) + y(x_1) & \text{if } x < x_1, \\ y(x) & \text{if } x_1 \le x \le x_n, \\ y'(x_n)(x - x_n) + y(x_n) & \text{if } x > x_n. \end{cases}$$

https://doi.org/10.1017/S1446181115000310 Published online by Cambridge University Press

4. Pricing by using Crank–Nicolson method

Once we have a computable local volatility function, we can use the finite-difference method to solve the PDE (3.1). Suppose that we would like to price a European call option with strike K and maturity T years. We approximate the PDE (3.1) with boundary conditions $V(S(T), T) = (S(T) - K)^+$.

4.1. Mesh properties First we need a mesh of (price, time) pairs. Suppose that we have N different time points and M different prices in the mesh. Furthermore, assume that the mesh is rectangular and uniformly spaced with boundaries of:

- 0 and T for the time axis;
- $S_0 \exp\{-D\}$ and $S_0 \exp\{D\}$ for the price axis, where $D = \gamma\bar{\theta}\sqrt{T}$ and $\bar{\theta}$ is the average of the at-the-money implied volatilities. We set $\gamma = 7$, which we determined experimentally as a value that resulted in an overall small numerical error for the Crank–Nicolson method. This value corresponds to a very small probability for the price to move beyond S_0e^D.

In addition, we scale the number of time points by T. Setting $N = 500T + 500$ gives sufficiently good results. Define $\Delta t = T/N$ and $\Delta x = 2D/M$. Then the time interval $[0, T]$ is discretized by

$$t_0 = 0, t_1 = \Delta t, t_2 = 2\Delta t, \ldots, t_N = T.$$

Then the price interval $[S_0e^{-D}, S_0e^D]$ is discretized by $s_i = S_0e^{-D+i(\Delta x)}$ for $i = 0, 1, \ldots, M$. This definition allows mesh points to coincide with the spot price. This is done for the convenience of the backtesting procedure to calculate difference vanillas using the same mesh. Next, the PDE (3.1) contains partial derivatives with respect to the logarithm of price. Let $x_i = \ln s_i$; then

$$x_{i+1} - x_i = \ln s_{i+1} - \ln s_i = \ln\left(\frac{S_0 \exp\{-D + (i+1)\Delta x\}}{S_0 \exp\{-D + i\Delta x\}}\right) = \Delta x,$$

so the price points are indeed uniformly spaced in terms of log-prices.

4.2. The finite-difference scheme Let $i \in \{1, \ldots, N\}$, $\nu(t, x) = r(t) - q(t) - [\sigma(e^x, t)]^2/2$ and $V_j^i = V(s_j, t_i)$. The *Crank–Nicolson* scheme is given by

$$a_j^{i-1}V_j^{i-1} - b_j^{i-1}V_{j+1}^{i-1} - c_j^{i-1}V_{j-1}^{i-1} = d_j^iV_j^i + b_j^iV_{j+1}^i + c_j^iV_{j-1}^i,$$

where

$$a_j^i = \frac{r(t_i)}{2} + \frac{1}{\Delta t} + \frac{\sigma^2(t_i, s_j)}{2(\Delta x)^2}, \qquad b_j^i = \frac{\sigma^2(t_i, s_j)}{4(\Delta x)^2} + \frac{\nu(t_i, s_j)}{4\Delta x},$$

$$c_j^i = \frac{\sigma^2(t_i, s_j)}{4(\Delta x)^2} - \frac{\nu(t_i, s_j)}{4\Delta x}, \qquad d_j^i = \frac{1}{\Delta t} - \frac{r(t_i)}{2} - \frac{\sigma^2(t_i, s_j)}{2(\Delta x)^2}.$$

For boundary conditions, we use the facts that

$$\lim_{S\to 0} \frac{\partial V}{\partial S} = e_0 = \begin{cases} 0 & \text{if } V \text{ is a call option,} \\ -1 & \text{if } V \text{ is a put option,} \end{cases}$$

https://doi.org/10.1017/S1446181115000310 Published online by Cambridge University Press

and

$$\lim_{S\to\infty} \frac{\partial V}{\partial S} = e_\infty = \begin{cases} 1 & \text{if } V \text{ is a call option,} \\ 0 & \text{if } V \text{ is a put option.} \end{cases}$$

This leads to the following equations:

$$V_0^i - V_1^i = e_0(s_0 - s_1), \quad V_M^i - V_{M-1}^i = e_\infty(s_M - s_{M-1}).$$

To initiate the scheme, we set for all $j = \{0, \ldots, M\}$

$$V_j^N = \begin{cases} (s_j - K)^+ & \text{if we are pricing a call option,} \\ (K - s_j)^+ & \text{if we are pricing a put option.} \end{cases}$$

We then repeatedly solve the system until we obtain $(V_1^0, V_2^0, \ldots, V_M^0)^{\mathrm{T}}$ (for details, see, for example, the book by Wilmott [13]). If M is an odd integer, the price of the option is $V^0_{(M+1)/2}$. Otherwise, we may fit an interpolating function $\hat{V}(s)$ through

$$(s_1, V_1^0), (s_2, V_2^0), \ldots (s_M, V_M^0)$$

and then the price is $\hat{V}(S_0)$. Also note that the delta of the option is $\hat{V}'(S_0)$; we have found that a natural cubic spline for $\hat{V}$ gives good results.

Remark 4.1. This pricing method is very fast if the mesh points of our local volatility function coincide with the mesh points in our finite-difference scheme. This is how we implemented our scheme; we first set the mesh points for our finite-difference scheme and then we pre-compute the local volatility function at these points.

5. Market data layout

In this paper, we work with daily AUD/USD implied volatility data from 22 March 2005 to 15 July 2011. For each trading day, the market data contains a spot price and, for a range of maturities (1 week, 1 month, 2 months, 3 months, 6 months, 1 year, 2 years, 3 years, 4 years and 5 years), there are:

- implied volatility for at-the-money (ATM) options;
- risk reversal for 10 and 25 delta calls, denoted by $RR_{10\Delta\mathrm{Call}}$ and $RR_{25\Delta\mathrm{Call}}$, respectively;
- butterfly for 10 and 25 delta puts, denoted by $Fly_{10\Delta\mathrm{Put}}$ and $Fly_{25\Delta\mathrm{Put}}$, respectively;
- zero rates (yields) for the domestic and foreign currencies.

From this data, we need to extract the strike prices and implied volatilities for traded vanilla options. This is done through the Black–Scholes framework. Taking the

Black–Scholes price of a call option (2.2) and differentiating, we obtain the *call delta*

$$\Delta_{\text{call}}(S_0, K, T, \gamma_d, \gamma_f, \sigma) = \frac{\partial C(S_0, K, T, \gamma_d, \gamma_f, \sigma)}{\partial S} = e^{-\gamma_f T}\Phi(d_1), \tag{5.1}$$

where $\gamma_d = (1/T)\int_0^T f(t)\,dt$ and $\gamma_f = (1/T)\int_0^T q(t)\,dt$ denote the domestic and foreign yields, respectively (this is discussed in detail in Section 5.3). Utilizing put–call parity, the *put delta*

$$\Delta_{\text{put}}(S_0, K, T, \gamma_d, \gamma_f, \sigma) = \Delta_{\text{call}}(S_0, K, T, \gamma_d, \gamma_f, \sigma) - e^{-\gamma_f T}. \tag{5.2}$$

In the following, we use 10ΔPut and 25ΔPut to denote the volatility σ and strike price K that give a put delta Δ_{put} of 10% and 25%, respectively. Similarly, 10ΔCall and 25ΔCall denote the volatility σ and strike price K that give a call delta Δ_{call} of 10% and 25%, respectively.

5.1. Computing implied volatilities With usual market definitions, we reconstruct the market implied volatilities by using the following formulae:

$$\sigma_{10\Delta\text{Put}} = \sigma_{\text{ATM}} + \text{Fly}_{10\Delta} - \tfrac{1}{2}RR_{10\Delta}, \quad \sigma_{25\Delta\text{Put}} = \sigma_{\text{ATM}} + \text{Fly}_{25\Delta} - \tfrac{1}{2}RR_{25\Delta},$$
$$\sigma_{25\Delta\text{Call}} = \sigma_{\text{ATM}} + \text{Fly}_{25\Delta} + \tfrac{1}{2}RR_{25\Delta}, \quad \sigma_{10\Delta\text{Call}} = \sigma_{\text{ATM}} + \text{Fly}_{10\Delta} + \tfrac{1}{2}RR_{10\Delta}.$$

5.2. Computing strikes After we determine the implied volatilities, the only parameter yet to be determined is the strike price. We use the delta formulae of (5.1) and (5.2) and an implementation of the inverse cumulative distribution function Φ^{-1} of the standard Normal distribution to determine the strike. For example, to obtain the strike price for the 10Δ*Put* option, we are looking for the value of K satisfying

$$\Delta_{\text{put}}(S_0, K, T, \gamma_d, \gamma_f, \sigma_{10\Delta\text{Put}}) = 0.1,$$

which is easy to calculate via the inverse normal distribution function.

5.3. Interest rates On each trading day, we can extract from the market the so-called *zero-coupon interest rates* or *zero rates* for a range of different maturities; we give an example to explain the meaning of these rates. Suppose we have the following market interest rate yields, 4.8%, 4.9%, 5% and 5.1% for maturities 1, 2, 3 and 4 years, respectively. With continuous compounding of interest rates, a one-year investment of AUD 10 grows to $10 \times e^{0.048\times 1} = 10.49$. A two-year investment of the same amount grows to $10 \times e^{0.049\times 2} = 11.03$.

We can now interpolate between these data points to obtain what is called a *zero curve*. There is no universally accepted way to perform this interpolation. Suppose that the market rates are given by $(T_1, \gamma_1), (T_2, \gamma_2), \dots, (T_n, \gamma_n)$, where T_i denotes the ith maturity and γ_i is its zero rate. We define our zero curve $\gamma(t)$ to be a function such that $t\gamma(t)$ is piecewise linear through the points

$$(0, \gamma_1), (T_1, \gamma_1), (T_2, \gamma_2), \dots, (T_n, \gamma_n).$$

Now we will need to have the instantaneous interest rates for various calculations, such as the Dupire formula (3.2). That is, we need to find the function $r(s)$ such that

https://doi.org/10.1017/S1446181115000310 Published online by Cambridge University Press

$\gamma(t) = (1/t)\int_0^t r(s)\,ds$. Since we assumed that $t\gamma(t)$ is piecewise linear, it implies that $r(s)$ is piecewise constant on the same intervals that $t\gamma(t)$ is piecewise linear.

By the construction, $\gamma(t) = \gamma_1$ for $t \in [0, T_1]$, which implies that $r(t) = \gamma(t)$ on $[0, T_1]$. Next, let $t \in (T_i, T_{i+1}]$ and consider

$$\gamma(T_{i+1})T_{i+1} - \gamma(T_i)T_i = \gamma_{i+1}T_{i+1} - \gamma_i T_i = \int_{T_i}^{T_{i+1}} r(s)\,ds = (T_{i+1} - T_i)r(t),$$

which yields

$$r(t) = \frac{\gamma_{i+1}T_{i+1} - \gamma_i T_i}{T_{i+1} - T_i} \quad \text{for } t \in (T_i, T_{i+1}].$$

In summary, the instantaneous interest rate is given by

$$r(t) = \begin{cases} \gamma_1 & \text{for } t \in [0, T_1], \\ \dfrac{\gamma_{i+1}T_{i+1} - \gamma_i T_i}{T_{i+1} - T_i} & \text{for } t \in (T_i, T_{i+1}],\ i = 1, \ldots, n-1. \end{cases} \tag{5.3}$$

5.4. Term structure of volatility for Black–Scholes To best compare the performance of Black–Scholes with local volatility, we need to have time-dependent volatilities for the Black–Scholes model. Under this condition, the market implied volatilities allow us to construct the term structure of volatility. Suppose that at-the-money implied volatilities are $(T_1, \sigma_1), (T_2, \sigma_2), \ldots, (T_n, \sigma_n)$. Then, similar to the instantaneous interest rate of the previous section, we define the instantaneous volatility $\sigma : [0, T_n] \to R^+$ by

$$\sigma(t) = \begin{cases} \sigma_1 & \text{for } t \in [0, T_1], \\ \sqrt{\dfrac{\sigma_{i+1}^2 T_{i+1} - \sigma_i^2 T_i}{T_{i+1} - T_i}} & \text{for } t \in (T_i, T_{i+1}],\ i = 1, \ldots, n-1. \end{cases}$$

From this, define $\sigma_{\text{avg}} : [0, T_n] \to R^+$ by $\sigma_{\text{avg}}(t) = \sqrt{(1/t)\int_0^t \sigma^2(s)\,ds}$. Then it is easy to check that $\sigma_i^2 = \sigma_{\text{avg}}^2(T_i)$ for all $i = 1, \ldots, n$.

6. Calibrating the model

We performed verification of local volatility model calibration procedures as follows.

(1) Obtain current market data and construct the local volatility surface as outlined in Section 3.2.

(2) For each market traded option V:

  (a) obtain V's strike price and maturity. Then apply the pricing methodology in Section 4 to obtain a price;

https://doi.org/10.1017/S1446181115000310 Published online by Cambridge University Press

(b) using the Black–Scholes formula (2.2), compute the implied volatility from the obtained price and compare with the market implied volatility of V. Ideally, the computed implied volatility should be the same as the market volatility but because of numerical errors we have a slight difference (see Section 7). We adjusted our procedures to obtain an absolute difference less than 0.5%.

Table 2 shows the average absolute difference for calibration errors for our implementation over each day of historical AUD/USD foreign exchange data.

7. Implementation

All our implementations were written in C++. Numerical errors from our implementation come from using a finite number of mesh points in the finite-difference method (Section 4.1) as well as from finite boundaries for the mesh.

Note that the pricing method for the local volatility model requires solving tridiagonal systems of equations for finding the natural cubic spline and for the finite-difference method. There exists an algorithm to solve the system in linear time (see the book by Press et al. [9, Section 2.4]).

8. Delta hedging

Let V denote the price of an option. The *delta* of the option is defined as $\Delta = \partial V/\partial S$. Delta is a measure of the sensitivity of the option price to changes in the value of the underlying asset. Under the Black–Scholes framework, Δ can be computed explicitly.

Suppose that we have a portfolio of options with stocks as underlying assets. *Delta hedging* is a strategy to reduce the risk of the portfolio to changes in price of the underlying assets. To hedge a short position of one call option we need to take a long position of Δ shares of the underlying asset. Because a change in share price leads to a change in delta, we must *rebalance* our long position to maintain the hedge. This means that if the current Δ changes to Δ', we must buy or sell to have Δ' shares. Under the Black–Scholes framework, the rebalancing must be performed continuously in time to obtain a riskless portfolio.

8.1. The delta hedging procedure Suppose that we are selling a European call option with expiry at time T, and that we wish to rebalance at N evenly spaced points in time. Let $\delta t = T/N$. Let $t_0 = 0, t_1 = \delta t, \ldots, t_N = T$. For $i \in \{1, \ldots, N\}$, let S_i and Δ_i denote the share price and delta of the call option at time t_i, respectively. Also, let the price of a call option at time t_0 be C. We note that calculation of C and Δ_i depends on the model we use for the asset price. Let $r(t)$ and $q(t)$ denote the instantaneous interest rate and instantaneous dividend yield at time t, respectively. These rates are computed by the formulae in Section 5.3. Then the delta hedging procedure is performed as follows.

(1) At t_0, we short one call for C cash, and go long Δ_0 shares. The cash position at this time is $P_0 = C - \Delta_0 S_0$.

(2) At t_1, we perform our first rebalancing. At this point of time, we need to have long Δ_1 shares, which results in a cash flow of $(\Delta_0 - \Delta_1)S_1$. To see why, suppose that $\Delta_0 < \Delta_1$. We need to buy $(\Delta_1 - \Delta_0)$ shares, which has a cash flow of $-(\Delta_1 - \Delta_0)S_1 = (\Delta_0 - \Delta_1)S_1$. On the other hand, if $\Delta_1 < \Delta_0$, we need to sell $\Delta_0 - \Delta_1$ shares, which has a cash flow of $(\Delta_0 - \Delta_1)S_1$.

(3) Next, note that interest charged (or accrued) on the cash position of $C - \Delta_0 S_0$ between times t_0 and t_1 is $(e^{r(t_0)\delta t} - 1)P_0$. Similarly, the continuous dividend yield paid/received over this time period is $(e^{q(t_0)\delta t} - 1)\Delta_0 S_0$.

(4) After rebalancing at t_1, our cash position is

$$P_1 = e^{r(t_0)\delta t} P_0 + (e^{q(t_0)\delta t} - 1)\Delta_0 S_0 + (\Delta_0 - \Delta_1)S_1.$$

(5) At t_2, we need to be long Δ_2 shares, which results in a cash flow of $(\Delta_1 - \Delta_2)S_2$. Again taking into account interest and dividend yield, our cash position at this time is

$$P_2 = e^{r(t_1)\delta t} P_1 + (e^{q(t_1)\delta t} - 1)\Delta_1 S_1 + (\Delta_1 - \Delta_2)S_2.$$

(6) In general, the cash position at time t_i, $i \in \{1, \ldots, N-1\}$, is

$$P_i = e^{r(t_{i-1})\delta t} P_{i-1} + (e^{q(t_{i-1})\delta t} - 1)\Delta_{i-1} S_{i-1} + (\Delta_{i-1} - \Delta_i)S_i.$$

(7) After rebalancing at time t_{N-1}, we have a cash position of P_{N-1} and a long position of Δ_{N-1} shares.

(8) At maturity $t_N = T$, we will sell our long position of shares. We still earn or pay interest and dividend over the period $[t_{N-1}, t_N]$. The final cash position is then

$$P_N = e^{r(t_{N-1})\delta t} P_{N-1} + (e^{q(t_{N-1})\delta t} - 1)\Delta_{N-1} S_{N-1} + \Delta_{N-1} S_N.$$

The hedging error is then defined as $P_N - (S_T - K)^+$.

8.2. Simulated delta hedging Under the Black–Scholes model, the asset price S follows geometric Brownian motion, where it is possible to have time-dependent drift and volatility, μ_t and σ_t, respectively. To simulate S, we use the scheme (see the book by Glasserman [7])

$$S(t_{n+1}) = S(t_n)\exp\{(\mu - \tfrac{1}{2}\sigma^2)\delta t + \sigma_{t_n}\sqrt{\delta t} Z_n\},$$

where Z_n are independent and identically distributed normal random variables with mean 0 and variance 1. Using this scheme to generate a trajectory of the price process S, we may then perform delta hedging.

Under the local volatility model, we simulate the asset price process S by

$$S(t_{n+1}) = S(t_n)\exp\{(\mu - \tfrac{1}{2}\sigma^2)\delta t + \sigma(S(t_n), t_n)\sqrt{\delta t} Z_n\}.$$

We performed simulated delta hedging under both the Black–Scholes and local volatility models and observed that hedging errors converged to zero as the time step decreases to zero.

https://doi.org/10.1017/S1446181115000310 Published online by Cambridge University Press

9. Historical delta hedging

In this section, we apply the delta hedging procedure with real daily AUD/USD implied volatility data as described in Section 5. For a given call option with maturity of T years, we set N to be the number of trading days between the day the option is written and the day of maturity. So, as in Section 8.1, we define $\delta t = T/N$ and $t_i = i\delta t$ for $i = 0, \ldots, N$. Then t_i represents the start of the $(i+1)$th trading day.

The instantaneous interest rates $r(t)$ and $q(t)$ represent the domestic USD and foreign AUD rates, respectively. The procedure for the historical backtest is the same as that described in Section 8.1. However, we must be careful with the interest rates, since for each trading day a new sequence of market zero rates is quoted. To be precise, the quantity $r(t_i)$ that is needed in the delta hedging procedure is obtained by taking the domestic zero rate for the nearest quoted maturity T_1 *from the market data corresponding to trading day* t_i (note the form of equation (5.3)).

The only points where the Black–Scholes and local volatility methods differ is the calculation of delta Δ_i on each trading day.

9.1. Backtesting under the Black–Scholes framework For each trading day t_i, define

$$\gamma_d^{(i)} = \frac{1}{T - i\delta t}\int_0^{T-i\delta t} r(t)\,dt, \quad \gamma_f^{(i)} = \frac{1}{T - i\delta t}\int_0^{T-i\delta t} q(t)\,dt,$$

$$\sigma_{avg}^{(i)} = \sqrt{\frac{1}{T - i\delta t}\int_0^{T-i\delta t} \sigma^2(t)\,dt},$$

where these quantities are obtained from the market data at time t_i. We note that the money implied volatilities are used.

At time t_0, we compute the initial call option price $C(S_0, K, T, \gamma_d^{(0)}, \gamma_f^{(0)}, \sigma_{\rm avg}^{(0)})$ by the Black–Scholes formula (2.2) and initial delta $\Delta_{{\rm call}(S_0,K,T,\gamma_d^{(0)},\gamma_f^{(0)},\sigma_{\rm avg}^{(0)})}$ given by equation (5.1). To obtain the cash position P_i at time t_i, we need the value of delta Δ_i at this time, which is computed as

$$\Delta_i = \Delta_{{\rm call}(S_i,K,T-i\delta t,\gamma_d^{(i)},\gamma_f^{(i)},\sigma_{\rm avg}^{(i)})}.$$

Note that we also used interpolated implied volatilities for each trading day instead of $\sigma_{\rm avg}^{(i)}$, and this gives us mostly indistinguishable results. Results for each method are shown in Tables 3 and 4.

9.2. Backtesting under the local volatility framework Under the local volatility model, we use the finite-difference scheme of Section 4.2 to compute the initial option price. Recall from Section 4.2 that the finite-difference scheme results in a sequence of prices $(V(s_1, t_0), V(s_2, t_0), \ldots, V(s_M, t_0))$ at time t_0, where s_i denotes the price grid points of the scheme. Fitting an interpolating function $\hat{V}(s)$ through

$$(s_1, V(s_1, t_0)), (s_2, V(s_2, t_0)), \ldots, (s_M, V(s_M, t_0)),$$

https://doi.org/10.1017/S1446181115000310 Published online by Cambridge University Press

TABLE 3. Mean and standard deviation of hedging errors for each model under consideration for one week maturity. Black–Scholes_1 denotes the Black–Scholes backtest using $\sigma_{\text{avg}}^{(i)}$ defined in Section 9.1 and Black–Scholes_2 denotes the method using interpolated implied volatilities.

Delta	Model	Mean	Std. dev.
10ΔPut	Black–Scholes_1	−0.0004	0.0024
10ΔPut	Black–Scholes_2	−0.0004	0.0025
10ΔPut	LocalVol_TC	−0.0001	0.0023
10ΔPut	LocalVol_TI	−0.001	0.0129
25ΔPut	Black–Scholes_1	−0.0004	0.0029
25ΔPut	Black–Scholes_2	−0.0004	0.0029
25ΔPut	LocalVol_TC	−0.0001	0.0028
25ΔPut	LocalVol_TI	−0.0009	0.0108
ATM	Black–Scholes_1	−0.0	0.0028
ATM	Black–Scholes_2	−0.0	0.0028
ATM	LocalVol_TC	−0.0003	0.0028
ATM	LocalVol_TI	−0.0008	0.0078
25ΔCall	Black–Scholes_1	0.0003	0.0023
25ΔCall	Black–Scholes_2	0.0002	0.0023
25ΔCall	LocalVol_TC	0.0002	0.0024
25ΔCall	LocalVol_TI	−0.0	0.0046
10ΔCall	Black–Scholes_1	0.0003	0.0016
10ΔCall	Black–Scholes_2	0.0003	0.0016
10ΔCall	LocalVol_TC	0.0003	0.0016
10ΔCall	LocalVol_TI	0.0003	0.002

the initial price is then given by $C = \hat{V}(S_0)$. We also define the time t_0 delta of the option by $\Delta_0 = \hat{V}'(S_0)$, the first derivative of V at S_0. Next, we describe two methods of computing the subsequent deltas, $\Delta_1, \dots, \Delta_{N-1}$. We first introduce some simplifying notions.

Let $i = 1, \dots, N-1$. Suppose that we apply the finite-difference scheme to the market data at time t_{-1}. After iteratively solving the required system of equations, we obtain a sequence of time t_{i-1} call option prices

$$(V(s_1, t_{i-1}), V(s_2, t_{i-1}), \dots, V(s_M, t_{i-1})).$$

We then define the function $\hat{V}_{(i-1)}(s)$ as the natural cubic spline passing through

$$(s_1, V(s_1, t_{i-1})), (s_2, V(s_2, t_{i-1})), \dots, (s_M, V(s_M, t_{i-1})).$$

9.2.1 *Theoretically correct delta.* The so-called *theoretically correct delta* Δ_i at time t_i for $i = 1, \dots, N-1$ is defined by $\Delta_i = \hat{V}'_{(i-1)}(S_i)$, where S_i is the spot price at time t_i.

https://doi.org/10.1017/S1446181115000310 Published online by Cambridge University Press

TABLE 4. Mean and standard deviation of hedging errors for each model under consideration for one month maturity.

Delta	Model	Mean	Std. dev.
10ΔPut	Black–Scholes_1	−0.0012	0.005
10ΔPut	Black–Scholes_2	−0.0012	0.005
10ΔPut	LocalVol_TC	−0.0006	0.004
10ΔPut	LocalVol_TI	−0.0063	0.0275
25ΔPut	Black–Scholes_1	−0.0008	0.0041
25ΔPut	Black–Scholes_2	−0.0009	0.0042
25ΔPut	LocalVol_TC	−0.0005	0.0039
25ΔPut	LocalVol_TI	−0.0051	0.0236
ATM	Black–Scholes_1	−0.0002	0.0036
ATM	Black–Scholes_2	−0.0003	0.0037
ATM	LocalVol_TC	−0.0006	0.0036
ATM	LocalVol_TI	−0.0035	0.0177
25ΔCall	Black–Scholes_1	0.0003	0.0029
25ΔCall	Black–Scholes_2	0.0003	0.0029
25ΔCall	LocalVol_TC	−0.0001	0.0029
25ΔCall	LocalVol_TI	−0.0016	0.0114
10ΔCall	Black–Scholes_1	0.0005	0.0018
10ΔCall	Black–Scholes_2	0.0004	0.0018
10ΔCall	LocalVol_TC	0.0003	0.002
10ΔCall	LocalVol_TI	−0.0005	0.0061

The idea behind this definition of delta is that if we compute a local volatility function from current market data with spot price S_0, a subsequent change in the spot price should not alter the local volatility function, that is,

$$\sigma(S, t; S_0) = \sigma(S, t; S_0 + \Delta S) \tag{9.1}$$

for some change in spot price ΔS. If the local volatility function fully captured the real diffusion process of the underlying asset, then (9.1) should prevail. However, there are claims in the literature that this is contrary to common market behaviour (see, for example, the article by Hagan et al. [8] and the book by Rebonato [10]).

9.2.2 *Sticky delta.* With *sticky delta*, we assume that a change in the spot price will not result in a change to the implied volatility and the delta [6]. That is, the market data implied volatilities (for example, in Table 1), which are expressed in terms of maturity and delta, do not change when the spot price changes. This leaves the strike price to be altered. If $\hat{\sigma}$ denotes the implied volatility, we can show that

$$\hat{\sigma}(K, T; S_0) = \hat{\sigma}\Big(K + \frac{K\Delta S}{S_0}, T; S_0 + \Delta S\Big).$$

https://doi.org/10.1017/S1446181115000310 Published online by Cambridge University Press

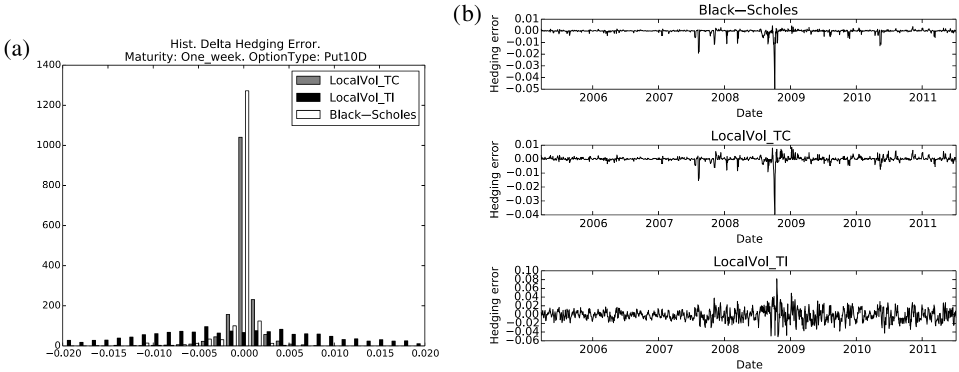

FIGURE 1. Delta hedging errors for (10ΔPut) calls with one week maturity. LocalVol_TC and LocalVol_TI denote results from the theoretically correct and theoretically incorrect (sticky delta) local volatility models, respectively. (a) Histogram. (b) Hedging error time series.

That is, under sticky delta a shift in the spot price S_0 by ΔS leads to a shifting of the market strike K by $K\Delta S/S_0$. To compute the quantity Δ_i under this assumption at time t_i, first compute the option price at time t_{i-1}, $\hat{V}_{(i-1)}(S_i)$. Then take the market data at time t_{i-1} and perturb the spot price S_{i-1} by a small quantity $\Delta S = (0.001)S_{i-1}$ (this quantity can neither be very small nor very large due to large errors introduced in the calculation of the derivative. Our chosen value for ΔS is determined from numerical tests for stability and accuracy). That is, define a new spot price $S^+_{i-1} = S_{i-1} + \Delta S$. Taking S^+_{i-1} as the new spot price and without modifying the implied volatilities, deltas and interest rates, recompute the market strike prices as explained in Section 5.2. Using this modified market data, compute a new option price by finite difference and interpolate through the t_{i-1} prices with the function $V^+_{(i-1)}(s)$. Similarly, define another spot price $S^-_{i-1} = S_{i-1} - \Delta S$, recompute a new set of strike prices, compute finite difference and interpolate through the resulting prices with the function $V^-_{(i-1)}(s)$. The *central difference sticky delta* is then defined as

$$\Delta_i = \frac{\hat{V}^+_{(i-1)}(S_i) - \hat{V}^-_{(i-1)}(S_i)}{2\Delta S}.$$

10. Results

Figures 1 to 10 depict histograms of delta hedging errors computed from the historical data under the frameworks of Black–Scholes and local volatility for different European calls. Recalling the notation of Section 5, a 10ΔPut call option denotes a European call option with volatility σ and strike price K such that $\Delta_{\text{put}=0.1}$. The 10ΔCall call option is defined similarly but instead with the condition $\Delta_{\text{call}=0.1}$. Within each histogram *LocalVol_TC* and *LocalVol_TI* denote the theoretically correct delta and sticky delta approaches, respectively. Sample means and standard deviations for these hedging errors are summarized in Tables 3 and 4.

https://doi.org/10.1017/S1446181115000310 Published online by Cambridge University Press

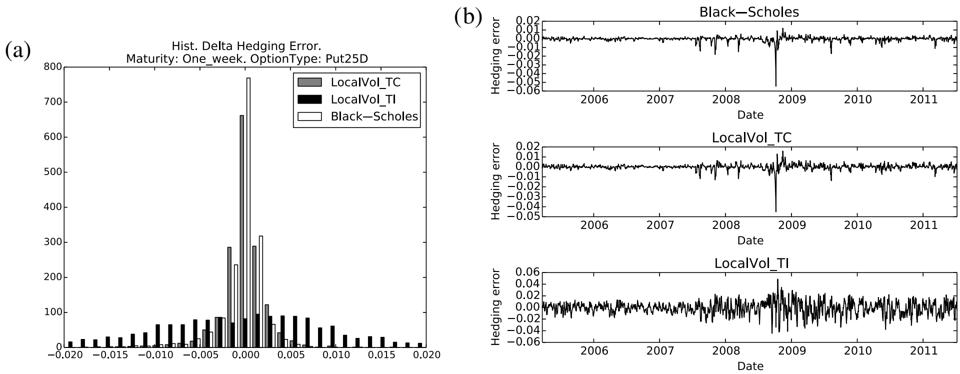

Figure 2. Delta hedging errors for (25ΔPut) calls with one week maturity. (a) Histogram. (b) Hedging error time series.

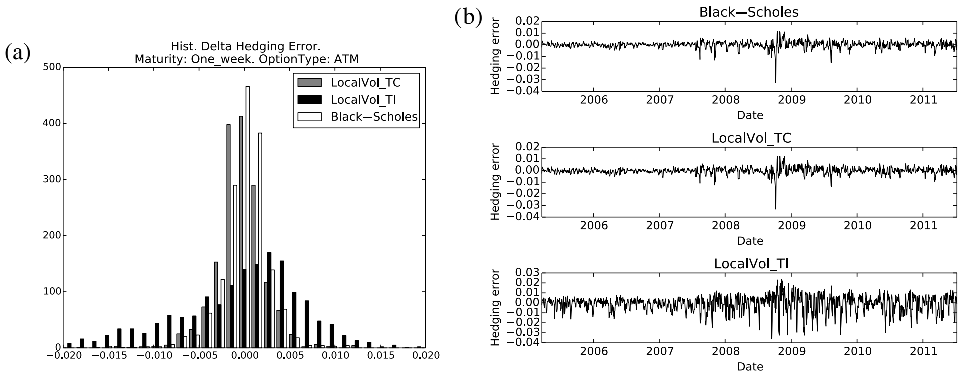

Figure 3. Delta hedging errors for ATM calls with one week maturity. (a) Histogram. (b) Hedging error time series.

We note that for in- and at-the-money options (Figures 1 to 3 for one week maturity and Figures 6 to 8 for one month maturity), the local volatility model with sticky delta performs significantly worse than the other two methods. It is only with deep in-the-money options (Figure 5a) that sticky delta local volatility exhibits hedging errors better than Black–Scholes.

11. Conclusion

Using delta hedging as the criterion to measure the effectiveness of a market model, our results show that Black–Scholes is no worse than the local volatility model. In fact, the Black–Scholes model performs significantly better than sticky delta local volatility, particularly for in- and at-the-money options. The theoretically correct delta local volatility model gives hedging errors which are not too far from those of Black–Scholes, and sticky delta local volatility performs noticeably worse

https://doi.org/10.1017/S1446181115000310 Published online by Cambridge University Press

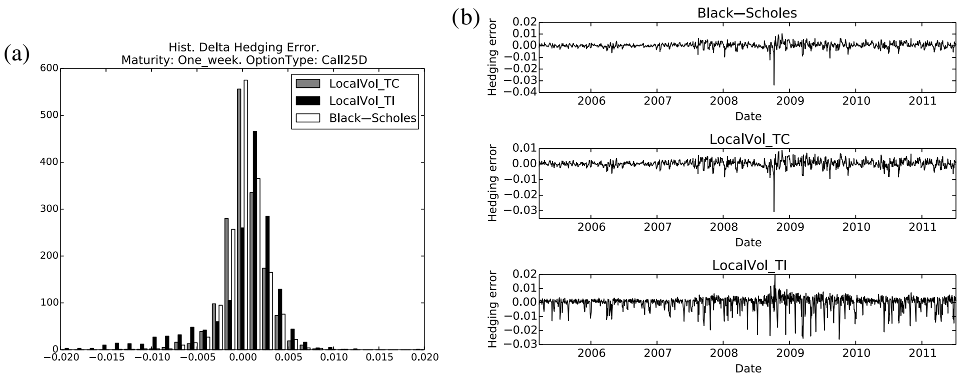

FIGURE 4. Delta hedging errors for (25ΔCall) calls with one week maturity. (a) Histogram. (b) Hedging error time series.

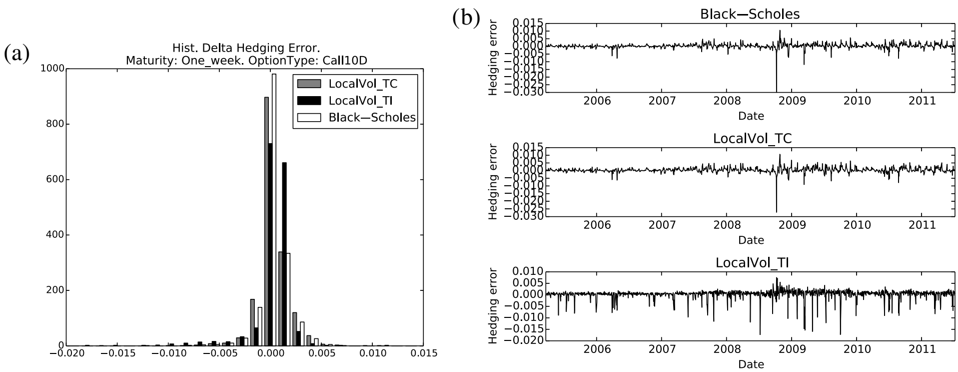

FIGURE 5. Delta hedging errors for (10ΔCall) calls with one week maturity. (a) Histogram. (b) Hedging error time series.

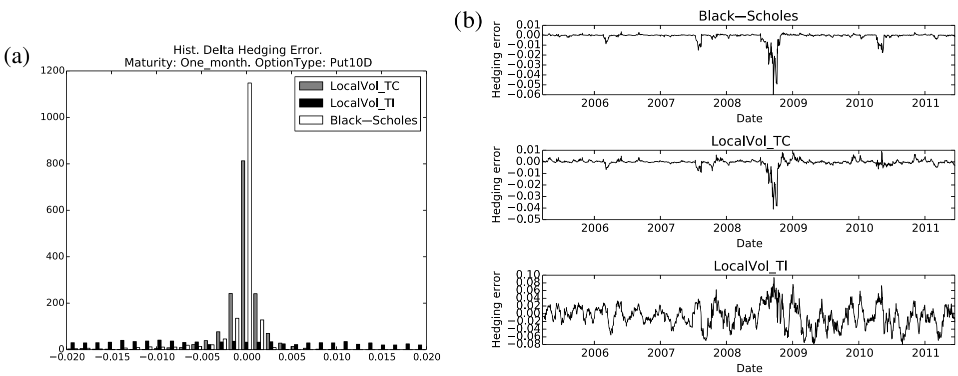

FIGURE 6. Delta hedging errors for (10ΔPut) calls with one month maturity. (a) Histogram. (b) Hedging error time series.

https://doi.org/10.1017/S1446181115000310 Published online by Cambridge University Press

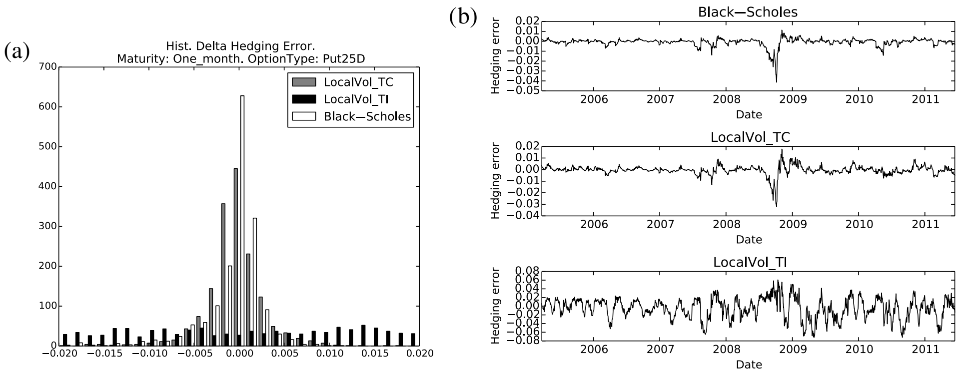

FIGURE 7. Delta hedging errors for (25ΔPut) calls with one month maturity. (a) Histogram. (b) Hedging error time series.

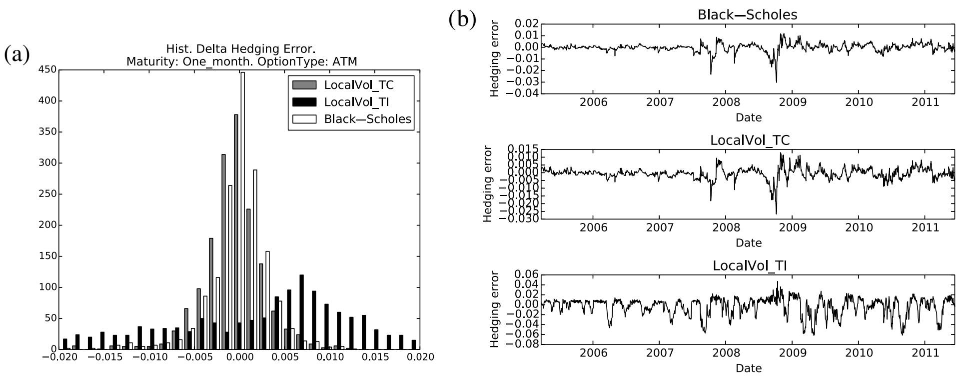

FIGURE 8. Delta hedging errors for ATM calls with one month maturity. (a) Histogram. (b) Hedging error time series.

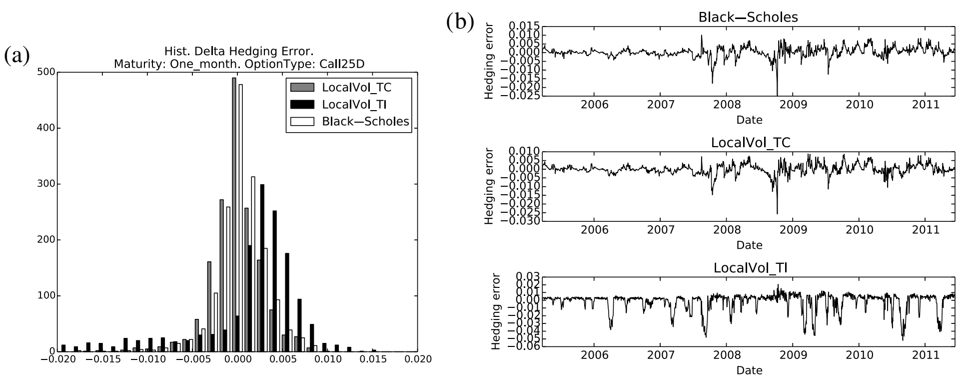

FIGURE 9. Delta hedging errors for (25ΔCall) calls with one month maturity. (a) Histogram. (b) Hedging error time series.

https://doi.org/10.1017/S1446181115000310 Published online by Cambridge University Press

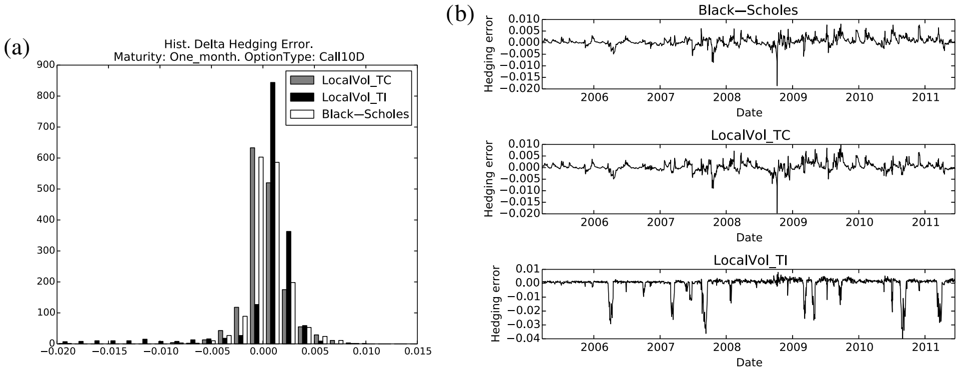

FIGURE 10. Delta hedging errors for (10ΔCall) calls with one month maturity. (a) Histogram. (b) Hedging error time series.

than the other models except for the case of deep out-of-the-money options. Further avenues of research include performing these empirical tests on other FX pairs and also incorporating other hedges. Also, the framework can be used to validate or compare other models, such as stochastic volatility or local stochastic volatility. It will also be of interest to determine hedging errors for exotic options, such as barrier options.

Acknowledgements

The authors would like to thank Igor Geninson and Rod Lewis from Commonwealth Bank of Australia Global Markets for their assistance as well as Xiaolin Luo from CSIRO. This first author T. G. Ling acknowledges the support of the Australian Research Council's Discovery Projects funding scheme (project number DP130103315).

References

[1] F. Black and M. Scholes, "The pricing of options and corporate liabilities", *J. Polit. Econ.* **81** (1973) 637–654; doi:10.1086/260062.

[2] G. Deelstra and G. Rayee, "Local volatility pricing models for long-dated FX derivatives", *Appl. Math. Finance* **20** (2013) 380–402; doi:10.1080/1350486X.2012.723516.

[3] E. Derman and I. Kani, "Riding on a smile", *Risk Mag.* **7** (1994) 32–39; http://www.math.ku.dk/kurser/2005-1/finmathtowork/DermanKaniRISK.PDF.

[4] B. Dupire, "Pricing with a smile", *Risk Mag.* **7** (1994) 18–20; http://www.math.ku.dk/kurser/2005-1/finmathtowork/DupireRISK.PDF.

[5] B. Feil, S. Kucherenko and N. Shah, "Volatility calibration using spline and high dimensional model representation models", *Wilmott J.* **1** (2009) 179–195; doi:10.1002/wilj.18.

[6] M. R. Fengler, *Semiparametric modelling of implied volatility* (Springer, Berlin, 2005).

[7] P. Glasserman, *Monte Carlo methods in financial engineering,* Volume 53 (Springer, New York, 2003); doi:10.1007/978-0-387-21617-1.

[8] P. S. Hagan, D. Kumar, A. S. Lesniewski and D. E. Woodward, "Managing smile risk", *Wilmott Mag.* (2002) 84–108; http://www.math.ku.dk/~rolf/SABR.pdf.

[9] W. H. Press, S. A. Teukolsky, W. T. Vetterling and B. P. Flannery, *Numerical recipes*, 3rd edn (Cambridge University Press, Cambridge, 2007).

[10] R. Rebonato, *Volatility and correlation: the perfect hedger and the fox*, 2nd edn (John Wiley & Sons Ltd, Chichester, 2004) 836; doi:10.1002/9781118673539.

[11] P. V. Shevchenko, "Advanced Monte Carlo methods for pricing European-style options", CMIS Technical Report CMIS 2001/148, CSIRO Mathematical and Information Sciences, 2001.

[12] R. White, "Local volatility", http://www.opengamma.com/sites/default/files/local-volatility-opengamma.pdf.

[13] P. Wilmott, *Paul Wilmott on quantitative finance*, 2nd edn (John Wiley & Sons Ltd, Chichester, 2006); http://www.amazon.com/Wilmott-Quantitative-Finance-Volume-Edition/dp/0470018704.