

Volatility Forecasting and the Efficiency of the Toronto 35 Index Options Market

Craig Doidge
Jason Z. Wei

University of Saskatchewan

Abstract

Existing research into Canadian options market efficiency is dated and focuses only on the stock options market. Until now no study has investigated the efficiency of the Toronto 35 index options market. This study fills the gap by examining the Toronto 35 index options market using a delta neutral straddle trading strategy. Straddle prices are computed by inserting different forecasts of volatility into the Black-Scholes option-pricing model. We consider the recently developed GARCH conditional-variance forecasts, implied-volatility forecasts, and forecasts that combine the information contained in historical return data and option prices in the market. Without transaction costs, the straddle strategy is able to generate trading profits when the same day's prices (or ex post prices) are used both to detect mispricing and to initiate a trade position. The profits disappear when transaction costs are introduced. When the straddle position is initiated based on the following day's prices (or ex ante prices), with or without transaction costs, no significant profits are available with any of the volatility forecasts. We therefore conclude that the Toronto 35 index options market is efficient with respect to the straddle trading strategy using different volatility forecasts.

Résumé

Les résultats des études antérieures portant sur l'efficacité du marché canadien des options ne sont déjà plus récents et ne concernent que le marché des options sur actions. L'efficacité relative aux options sur indice n'a encore été l'objet d'aucune recherche publiée. Notre étude vient combler ce vide, en analysant le marché des options sur l'indice TSE-35, à l'aide de stratégies de positions doubles à delta neutre. Les prix des positions doubles sont calculés en insérant différentes prévisions de volatilité dans le modèle de Black et Scholes. Les prévisions de volatilité que nous employons sont de trois types : (1) conditionnelles de type GARCH, (2) implicites et (3) dérivées d'information mixte combinant les rendements historiques et les valeurs marchandes des options. La stratégie de positions doubles génère des profits lorsque les prix quotidiens observés (les prix ex-post) sont utilisés à la fois pour détecter les écarts de prix et pour construire la stratégie de placement. Par contre, les profits disparaissent lorsque les frais de transactions sont pris en compte. Lorsque la stratégie de positions doubles repose sur les prix du jour à venir (les prix ex-ante), aucun profit significatif ne peut être généré, peu importe la méthode de prévision de la volatilité. De plus, ce résultat est valide, avec ou sans frais de transaction. Nous en concluons que le marché des options sur l'indice TSE-35 est efficient par rapport à une stratégie de positions doubles et aux méthodes de prévision utilisées dans cette recherche.

The concept of capital market efficiency is central to the field of finance; as such, financial economists have studied the efficiency of capital markets for many years. The concept of efficiency, as it applies to information available in the marketplace, is important because it

The authors are grateful to Professor Guy Bellemere for his authoritative translation of the abstract, and also two anonymous referees for useful comments.

Address all correspondence to Jason Wei, College of Commerce, University of Saskatchewan, 25 Campus Drive, Saskatoon, SK, Canada, S7N 5A7.

implies that in an efficient market, prices fully reflect all available information. That is, there are no exploitable correlations between future market returns and current information; investors who employ active trading rules cannot outperform the average market.

Black and Scholes (1972) and Galai (1977,1978) were the first researchers to investigate options market efficiency. Since then, several authors have investigated the efficiency of a number of different options markets around the world, particularly in the United States. There are two different research approaches: boundary-condi-

tions tests (Bhattacharya, 1983; Galai, 1978; Halpern & Turnbull, 1985), and some type of model-based test (Chance, 1986; Galai, 1977; Harvey & Whaley, 1992; Lamoureux & Lastrapes, 1993; Noh, Engle, & Kane, 1994; Stein, 1989). The former does not require any modeling assumptions, but leads to only a loose test in that it cannot detect small price deviations. The latter leads to much more powerful tests, but is subject to model-specification errors.

In Canada several studies of options market efficiency have been conducted. For example, Chua and Mokkelest (1988), Halpern and Turnbull (1985), and Mandron (1988) studied the efficiency of the Canadian stock options market in the late 1970s and early 1980s. They found strong evidence of inefficiency, which they attributed to thin trading, nonsimultaneity of prices, unfamiliarity of the market, and a noncompetitive trading environment. However, these studies are rather dated and only focus on the stock options market. To date, there is no evidence regarding the efficiency of the Toronto 35 index options market, an important index options market in Canada.

In this article, we use model-based tests to investigate the efficiency of the Toronto 35 index options market. A delta neutral straddle trading strategy driven by volatility forecasts is used. Employing volatility forecasts to study options-market efficiency is advantageous because volatility is a key determinant of option prices. In addition, because volatility can be computed from either historical returns or option prices, two distinct sources of information can be exploited. The article attempts to fill a major gap in the literature with respect to understanding Canadian options markets.

The remainder of this article is organized as follows. First we describe the data sources and the steps taken to avoid some potential data-related bias problems. Then we present the methods used to forecast volatility and examine the statistical accuracy of the forecasts. Third, we discuss the straddle tests and other issues related to market efficiency. This is followed by the major empirical results, and finally the conclusion.

Data

In May 1987, the Toronto Stock Exchange (TSE) introduced the Toronto 35 index, consisting of the 35 largest and most liquid stocks in Canada, to facilitate index-based trading and the trading of derivative products. Data for Toronto 35 index options (TXO options hereinafter) from August 1988 through July 1995 were obtained from the TSE. The daily record for each option contract contains the closing index level, exercise price, expiration date, last bid-ask prices, last price, and num-

ber of contracts traded. TXO options are European in nature with strike prices at increments of 5 points.¹ Typically, expiration months are the 3 consecutive near months, but maturities can be as long as 12 months. Expiration occurs on the third Friday of the expiration month and the last trading day is Thursday before the third Friday. All contracts are settled in cash.

Besides the option prices, measures of the risk-free interest rate and expected dividend yield on the index are required. The average of the bid-ask rates of the Treasury bill that expires closest to the option expiration date is used as a proxy for the risk-free interest rate. Each day, the annual dividend yield of the Toronto 35 index is used as a proxy for the annual expected dividend yield. The dividend yield is obtained from the daily newspaper, *The Globe and Mail*.

To mitigate potential biases, option quotes with daily volume of less than 10 contracts, closing prices lower than \$0.10, and time-to-maturity less than 7 days are excluded from the sample. In addition, quotes violating lower boundary conditions are eliminated from the sample. These filters serve to remove the contracts that cause most of the bias problems and focus the analysis on the near-term, close-to-the-money contracts. A summary of the filtered data set is given in Table 1.

Table 1 shows that 47.72% of the contracts in the sample have less than 31 days to expiration; 34.14% of the contracts have maturities between 31 and 60 days, and the remaining 18.14% have maturities of more than 60 days. The near-term expiration contracts account for 65.84% of total trading volume. With respect to moneyness, 58.43% of the contracts are close-to-the-money; 9.02% of the contracts are in-the-money, and 32.55% of the contracts are out-of-the-money. The close-to-the-money contracts account for 73.02% of total trading volume.

Volatility Estimation and Forecasting

To implement our straddle trading strategy, we need forecasts of the volatility of Toronto 35 index returns. The most common approaches are computing time-series forecasts based on historical-return data and computing the volatility implied in observed option prices. Another, less common, approach is to combine time-series forecasts with implied volatility. We will explore all three approaches in this study.

Historical return-based volatility forecasting. A simple method of forecasting volatility is to compute the variance of returns realized over the recent past and use it as an estimate of future volatility. Specifically, it is calculated as,

$$\sigma_t^2 = \bar{\sigma}^2 = \frac{1}{N-1} \sum_{i=1}^N (R_i - \bar{R})^2, \quad (1)$$

where \bar{R} is the mean daily Toronto 35 index return and N

Table 1
Summary of Options Data Set^a

	Total (#)	Percentage of total	Calls (#)	Percentage of calls	Puts (#)	Percentage of puts
Time-to-maturity						
Near-term ^b	12,312	47.72	6,005	49.62	6,307	46.04
Mid-term ^c	1,523,146	65.84	761,327	68.72	761,819	63.19
Long-term ^d	8,809	34.14	4,056	33.52	4,753	34.70
	584,477	25.26	253,579	22.89	330,898	27.45
	4,680	18.14	2,041	16.86	2,639	19.26
	205,826	8.90	92,984	8.39	112,842	9.36
Moneyness						
Close-to-the-money ^e	15,076	58.43	7,211	59.59	7,865	57.41
In-the-money	1,689,209	73.02	813,492	73.43	875,717	72.64
Out-of-the-money	2,327	9.02	961	7.94	1,366	9.97
	94,771	4.10	39,601	3.57	55,170	4.58
	8,398	32.55	3,930	32.47	4,468	32.62
	529,469	22.89	254,797	23.00	274,672	22.78

Note. Numbers in each cell are the number of contracts and trading volume, respectively. There are a total of 25,801 contracts in the data set; total trading volume is 2,313,449 contracts. The top half of the table shows the number of contracts and trading volume within each expiration class.

^aThere are a total of 1,743 trading days in the sample period from August 1988 to July 1995.

^bNear-term contracts are defined as options that have maturities of 30 days or less.

^cMid-term options have maturities between 31 days and 60 days.

^dLong-term options have maturities of 61 days or more.

^eClose-to-the money contracts are defined as options whose strike price is no more than 5 points away from the index level.

is the total number of observations. The past 100 trading days are used to compute the historical variance.

The GARCH model, developed by Bollerslev (1986) and Engle (1982), has become very popular in the finance literature because of its ability to capture the stochastic nature of volatility.² Further, it is generally known that a parsimonious GARCH(1, 1) specification is sufficient to capture the main dynamic of stock return series. The GARCH(1, 1) model involves the joint estimation of a conditional-mean equation and a conditional-variance equation, as follows:

$$R_t = \Phi_0 + \Phi_1 R_{t-1} + \epsilon_t \quad (2)$$

$$\sigma_t^2 = \alpha_0 + \alpha_1 \epsilon_{t-1}^2 + \beta_1 \sigma_{t-1}^2 \quad (3)$$

$$\epsilon_t | \Psi_{t-1} \sim N(0, \sigma_t^2) \quad (4)$$

where R_t is the Toronto 35 index return at time t , ϵ_t is the error term, σ_t^2 is the conditional variance of R_t , Ψ_{t-1} is the information set of all information up to time $t-1$, and

$N(0, \sigma_t^2)$ denotes a normal distribution with mean zero and variance σ_t^2 . To ensure a well-defined process, it is required that $(\alpha_0, \alpha_1, \beta_1) \geq 0$ and $\alpha_1 + \beta_1 < 1$. The lagged index return is included in Equation (2) to remove the serial correlation induced by nonsynchronous trading of comprising stocks in the underlying index.

Assuming a conditionally normal error distribution, the GARCH model is estimated using the method of maximum likelihood and the algorithm developed by Berndt, Hall, Hall, and Hausman (1974).³ The model's unconditional variance is used as the initial conditional variance in estimation. The period from August 2, 1988 to December 31, 1991 is used to obtain initial estimates of the model's parameters. Once the initial parameters are estimated, a rolling sample of 600 observations is used to update the parameters on a daily basis. The updated parameters are then used to compute new volatility forecasts.

One-period-ahead volatility estimates can be calculated directly from the explicit specification of the model's conditional variance. Following Heynen,

Kemna, and Vorst (1994), the time t forecast of the average variance of returns over the next n periods to time T is computed as,

$$\sigma_{i,T}^2 = \frac{1}{n} \sum_{i=1}^n E_t[\sigma_{i,t+i}^2] = \bar{\sigma}^2 + (\sigma_{i,t+1}^2 - \bar{\sigma}^2) \cdot \frac{1 - \kappa^{(T-t)}}{(T-t)(1-\kappa)} \quad (5)$$

where $\kappa = \alpha_1 + \beta_1$ and $\bar{\sigma}^2$ is the model's unconditional variance, which equals $\alpha_0 / (1 - \kappa)$.⁴

Although the GARCH(1, 1) model seems to work well for modeling stock returns, the simple structure of the model forces it to be symmetric in its response to past errors. However, Black (1976) and Christie (1982) found that bigger negative errors tend to be associated with higher future volatilities, which is believed to be caused by the firm's leverage position. To capture the leverage effect, Nelson's (1991) EGARCH model is employed.⁵ In this model, Equation (3) of the GARCH model is changed to

$$\ln \sigma_t^2 = \alpha_0 + \beta_1 \ln \sigma_{t-1}^2 + \beta_2 \xi_{t-1} + \beta_3 \left(|\xi_{t-1}| - \sqrt{\frac{2}{\pi}} \right) \quad (6)$$

where ξ_t , the standardized residual is a standard normal variate. A one-day-ahead volatility estimate can be obtained directly from the explicit specification of the model's conditional variance. According to Heynen et al. (1994), the more-than-one-day-ahead forecast (from time t to time T) can be calculated as where

$$\bar{\sigma}_{i,T}^2 = \frac{\sigma^2}{C(\beta_1, \beta_2, \beta_3)} \sum_{i=1}^n \sigma_{i+1}^{2\beta_1^{i-1}} \cdot \exp \left(\frac{- \left(\alpha_0 - \beta_3 \sqrt{\frac{2}{\pi}} \right) \beta_1^{i-1}}{1 - \beta_1} \right) \cdot \exp \left(\frac{\frac{1}{2} (\beta_2^2 + \beta_3^2) \beta_1^{2(i-1)}}{1 - \beta_1^2} \right) \cdot C_n(\beta_1, \beta_2, \beta_3) \quad (7)$$

$$\bar{\sigma}^2 = \exp \left(\frac{- \alpha_0 - \beta_3 \sqrt{\frac{2}{\pi}}}{1 - \beta_1} + \frac{1}{2} \frac{(\beta_2^2 + \beta_3^2)}{1 - \beta_1^2} \right) \cdot C(\beta_1, \beta_2, \beta_3),$$

$$C(\beta_1, \beta_2, \beta_3) = \prod_{m=0}^{\infty} [F_m(\beta_1, \beta_2, \beta_3) + F_m(\beta_1, -\beta_2, \beta_3)],$$

$$F_m(\beta_1, \beta_2, \beta_3) = N[\beta_1^m (\beta_3 + \beta_2)] \cdot \exp[\beta_1^{2m} \beta_2 \beta_3],$$

$$C_1 = 1 \text{ and } C_i = \prod_{m=0}^{i-2} [F_m(\beta_1, \beta_2, \beta_3) + F_m(\beta_1, -\beta_2, \beta_3)] \text{ for } i \geq 2.$$

The EGARCH model parameters are estimated using the same algorithm as in GARCH(1, 1) estimation.

Implied volatility forecasts. Another approach to forecasting volatility is to compute the volatility implied in option prices. We use the Black and Scholes (1973) model to accomplish this. The European call-and-put option prices are given by

$$C = S \cdot e^{-\gamma \cdot (T-t)} \cdot N(d_1) - X \cdot e^{-r \cdot (T-t)} \cdot N(d_2), \quad (8)$$

$$P = X \cdot e^{-r \cdot (T-t)} \cdot N(-d_2) - S \cdot e^{-\gamma \cdot (T-t)} \cdot N(-d_1), \quad (9)$$

where $d_1 = [\ln(S/X) + (r - \gamma + \sigma^2/2)(T-t)] / \sigma \sqrt{(T-t)}$, $d_2 = d_1 - \sigma \sqrt{(T-t)}$, S is the index level, X is the exercise price, r is the riskfree rate, γ is the annual dividend yield, and $(T-t)$ is the time remaining to maturity. The only unknown input is σ . But once an option's price is given we can invert the formula and obtain an implied volatility, which can in turn be used as a forecast.

In efficient markets, all options on the same underlying asset should imply the same level of volatility if the Black-Scholes model is the correct pricing model. However, it is well known that different options on the same underlying asset often imply different levels of volatility at a given point in time. Day and Lewis (1988) suggest that this result can be explained by noise induced by nonsynchronous quotes and the bid-ask spread. They and Harvey and Whaley (1992) argue that to decrease the measurement errors induced by these biases, the volatility measure must incorporate information from different options that are traded on the same day.

To compute a measure of implied volatility each day, the generalized least squares (GLS) method (Day & Lewis, 1988) is used. An advantage of this method is that it allows an explicit weighting scheme to reduce prediction errors. Because the GLS approach weighs each observation in proportion to the day's total trading volume, the implied volatility estimate tends to be influenced by near-term, close-to-the-money contracts, which is a desirable result.⁶

The GLS scheme is usually applied separately to calls, puts, and each expiration series. However, due to the relatively low trading volume of Toronto 35 index options, the estimation is implemented in a slightly different manner in this study. Specifically, all options (calls and puts) of all maturities are included in the weighting scheme each day. To illustrate, let $O_j(\sigma_0)$ denote the theoretical price of option j , given the volatility estimate σ_0 . The actual price of option j is denoted by O_j . The price discrepancy is defined as,

$$Y_j = O_j - O_j(\sigma_0). \quad (10)$$

If J is the total number of option prices on a given day, then given an initial estimate of volatility σ_0 , the new estimate of the implied volatility $\sigma_{i,T}$ is chosen to minimize,

$$\sum_{j=1}^J [\omega_j Y_j]^2, \tag{11}$$

where ω_j is the proportion of trading volume of option j relative to the total trading volume of that day. At each iteration of this process, the new estimate of the implied volatility is given by,

$$\sigma_{i,T} = \sigma_{(0)} + \sum_{j=1}^J \omega_j^2 V_j Y_j / \sum_{j=1}^J \omega_j^2 V_j^2 \tag{12}$$

where V_j is the partial derivative of option j 's price with respect to the underlying volatility. When the estimate converges to within a given tolerance level, the estimate $\sigma_{i,T}$ is taken as acceptable. If the estimate is not acceptable, the procedure is repeated, using $\sigma_{i,T}$ in place of σ_0 . The final estimate of implied volatility following the GLS scheme represents the average volatility expected over the remaining life of the option under consideration.⁷

Combination forecasts. Clemen (1989) states that the idea of combining forecasts implicitly assumes that one cannot correctly identify the underlying process followed by the variable of interest. Previous research indicates that neither implied volatility nor GARCH conditional-variance forecasts contain all relevant information about future volatility. Therefore, following Vasilellis and Meade (1996), we combine information from the stock market with information from the options market to predict future volatility. While combinations can be formed from any number of different forecasts, only two different forecasts are combined in this study—a GARCH conditional variance and the implied volatility.

Three combination schemes are considered. The simplest combination scheme is to calculate a simple average that assigns an equal weight to each forecast:

$$W_{G,t} = W_{IV,t} = 0.50, \tag{13}$$

where $W_{G,t}$ and $W_{IV,t}$ are the weights given to the GARCH and implied volatility forecasts. The combination forecast is then computed as

$$\sigma_{i,T}^{(C)} = W_{G,t} \cdot \sigma_{i,T}^{(G)} + W_{IV,t} \cdot \sigma_{i,T}^{(IV)} \tag{14}$$

where $\sigma_{i,T}^{(G)}$ and $\sigma_{i,T}^{(IV)}$ denote the GARCH and implied volatility forecasts, respectively. The second weighting scheme refines the simple average by allowing the forecast that has performed best in the recent past to receive a heavier weight. To this end, a forecast error is computed (for each individual forecast) by subtracting the volatility forecast from the actual volatility realized over the forecast period. Given the GARCH forecast and the implied volatility forecast, the appropriate weight for the GARCH forecast is

$$W_{G,t} = \left(\sum_{i=t-T-P}^{t-T-1} err_{G,i}^2 \right)^{-1} / \left[\left(\sum_{i=t-T-P}^{t-T-1} err_{G,i}^2 \right)^{-1} + \left(\sum_{i=t-T-P}^{t-T-1} err_{IV,i}^2 \right)^{-1} \right], \tag{15}$$

where err_G is the GARCH forecast error, err_{IV} is the implied volatility forecast error, P is set equal to 40 trading days, and T is the forecast horizon. T is included in the equation to ensure that all forecast errors are known at time t . The appropriate weight for the implied volatility forecast is computed analogously. Equation (14) is again used to compute the combined forecast. As shown towards the end of this section, GARCH(1, 1) performs slightly better than EGARCH(1, 1) in forecasting future volatilities, so we will use forecasts from GARCH(1, 1) to calculate the average forecasts described above.

Finally, following Day and Lewis (1992) and Lamoureux and Lastrapes (1993), the third combination scheme allows the GARCH model to incorporate forward-looking information in estimation. Specifically, Equation (3) of the GARCH model is modified such that the daily estimate of implied volatility is added as an exogenous variable to the variance specification:

$$\sigma_t^2 = \alpha_0 + \alpha_1 \epsilon_{t-1}^2 + \beta_1 \sigma_{t-1}^2 + \zeta_1 \sigma_{t-1}^{2(IV)} \tag{16}$$

where ζ_1 is the coefficient of the lagged implied volatility parameter. The standard GARCH model discussed earlier is then reestimated.⁸ For brevity, we call this model GARCH/IV model. Forecasts of future volatility with the GARCH/IV model are computed by,

$$\sigma_{t+n}^2 = \alpha_0 + \alpha_1 E[\epsilon_{t+n-1}^2 | \psi_t] + \beta_1 \sigma_{t+n-1}^2 + \zeta_1 E[\sigma_{t+n-1}^{2(IV)} | \psi_t]. \tag{17}$$

To evaluate the second expectation term, it is assumed that implied volatility is constant over the forecasting period.⁹ The following recursive substitution procedure is then employed to compute volatility forecasts:

$$\sigma_{t+n}^2 = \alpha_0 + \alpha_1 \sigma_{t+n-1}^2 + \beta_1 \sigma_{t+n-1}^2 + \zeta_1 \sigma_t^{2(IV)}. \tag{18}$$

To recapitulate, we have proposed three historical return-based forecasts, one implied volatility forecast, and three combination forecasts. The historical return-based forecasts are the simple historical variance, the GARCH(1, 1) forecast, and the EGARCH(1, 1) forecast. The combination forecasts are the simple average, the weighted average (of the implied and the GARCH(1, 1) forecasts), and the GARCH / IV forecast.

Testing Methodology

Table 2
Statistical Evaluation of Volatility Forecasts

Forecast	MAE ^a	MAPE (%) ^b	RMSE ^c
Historical volatility	0.0246	20.59	0.0304
GARCH(1,1)	0.0222	18.40	0.0287
EGARCH(1,1)	0.0223	18.58	0.0289
Implied volatility	<i>0.0276</i>	<i>24.71</i>	<i>0.0337</i>
Simple average ^d	0.0232	20.29	0.0289
Weighed average ^d	0.0217	18.80	0.0272
GARCH/IV	0.0207	17.06	0.0274

Note. Italicized numbers are the largest number in the column, and numbers in bold type are the smallest.

^aMAE: Mean absolute error

^bMAPE: Mean absolute percent error

^cRMSE: Root mean square error

^dBased on the historical volatility and the GARCH(1, 1) volatility.

The straddle tests used to test market efficiency depend critically on our forecasts of volatility. Therefore, it is important to check the accuracy of the forecasts before we attempt to use them to trade straddles. Statistical tests are employed to evaluate the performance of the forecasting techniques over a 1-month forecast horizon. The first part of the sample period, August 1988 to December 1990 is used to estimate the initial parameters of the time-series models. The period from January 1991 to July 1995 (1,137 trading days) is used on a rolling basis to evaluate out-of-sample forecasting performance. The out-of-sample forecasts are compared to the actual volatility realized over the forecast horizon. The mean absolute error (MAE), mean absolute percent error (MAPE) and the root mean square error (RMSE) statistics are calculated. The results are summarized in Table 2.

It can be seen that the GARCH/IV combination forecast has the lowest MAE and MAPE statistics, at 0.0207 and 17.06%, respectively. The implied volatility forecast has the highest MAE and MAPE, at 0.0276 and 24.71%, respectively. For the RMSE, the weighted average combination forecast is the lowest at 0.0272, and the implied volatility is the highest at 0.0337. The combination forecasts seem to be the most accurate, followed by the GARCH forecasts, historical volatility, and implied volatility. Between the two GARCH forecasts, GARCH(1, 1) outperforms EGARCH(1, 1) by a slight margin. Overall, the forecasts appear to be reasonably accurate; the statistics are in line with similar results reported by Akgiray (1989), Brailsford and Faff (1996), and Vasilellis and Meade (1996).

Two issues must be resolved before we discuss the trading strategy used to test market efficiency. The first has to do with the distinction between ex post and ex ante tests. In an ex post test, it is assumed that a particular trading strategy can be executed at the same price as that used to identify the profitable trade in the first place. An ex post test therefore cannot conclude the status of market efficiency because the transactions are not realistic. Nonetheless, it is a first step in detecting the true state of the market. In contrast, an ex ante test assumes that a trading strategy will be executed at the next set of available prices after the profitable trade has been identified. For example, if a price deviation is detected using today's closing prices and a trading strategy is designed accordingly, then it is more realistic to assume that the strategy will be implemented at tomorrow's prices. It is clear that an ex ante test is indeed a more direct test of market efficiency. For completeness, we will perform both types of test.

The second issue concerns transaction costs. There are two components of transaction costs: commission fees and the bid-ask spread, both of which will be considered in this study. The incorporation of commission fees is straightforward. Considering the average option prices in the sample and industry practices, we assume a commission fee of \$0.025 per option.¹⁰ When including a commission fee, we assume that transactions can be settled at the closing prices. To incorporate the bid-ask spread as a form of transaction cost, we assume that options are bought at the ask prices and sold at the bid prices. It should be noted that this assumption tends to exaggerate the true costs, because some transactions do get settled at prices within the bid-ask range. The exaggerated costs tend to cause the tests to be biased in favour of the efficient market hypothesis.

Because a distinction made between ex post and ex ante tests and the cost of the bid-ask spread may be important, six separate tests are conducted: (a) ex post closing price test, (b) ex post closing price test with transaction costs, (c) ex post bid-ask price test, (d) ex ante closing price test, (e) ex ante closing price test with transaction costs, and (f) ex ante bid-ask price test.

The trading strategy employed is similar to that of Noh, Engle, and Kane (1994), who traded at (near)-the-money straddles rather than individual options.¹¹ This strategy is attractive for a number of reasons. Trading straddles does not involve transacting in any other type of security that may trade on another exchange. This simplifies any assumptions that must be made about transaction costs; it also avoids difficulties associated with lags in information flow and in the execution of orders placed on other exchanges. Furthermore, because

Table 3
Straddle Prices and Buy/Sell Decisions

Forecast	Average call price ^a	Average put price ^b	Average straddle price ^c	Buy decisions (#)	Sell decisions (#)
Historical volatility	2.47	2.15	4.62	277	790
GARCH(1,1)	2.41	2.10	4.51	263	804
EGARCH(1,1)	2.41	2.09	4.50	267	800
Implied volatility	2.69	2.37	5.06	605	462
Simple average	2.55	2.23	4.78	354	713
Weighed average	2.52	2.20	4.72	331	736
GARCH/IV	2.38	2.06	4.44	115	952

Notes: There are 1,067 trading days in the test period from January 1991 to July 1995.

^aThe average closing market call option price over the sample period is \$2.65.

^bThe average closing market put option price is \$2.35

^cThe average closing market straddle price over the sample period is \$5.00.

at-the-money straddles are delta neutral, there is no need to delta-hedge the position.

The straddle data set is constructed from the set of options that met the selection criteria discussed in the data section. From this set, the nearest-to-the-money matching call and put options with the greatest trading volumes are used to construct straddle positions. There are 1,137 potential trading days in the sample period from January 1991 to July 1995; however, only 1,067 are used because it was not always possible to initiate an appropriate straddle position (i.e. an at-the-money straddle could not be constructed for certain days).

To identify mispriced straddles, the Black-Scholes model is used. When closing option prices are used, if the straddle is underpriced (overpriced) relative to the model price, it is selected for purchase (sale). The position is reversed at the next set of available closing prices. When bid-ask prices are used, if the closing straddle price is less (greater) than the model price, the straddle is purchased (sold) at the ask (bid) prices. The purchase (sale) transaction is reversed at the next available set of bid (ask) prices.¹²

Black-Scholes straddle prices are computed daily, given the relevant volatility forecast under consideration. All volatility forecasts used in the Black-Scholes model are specified as the average volatility expected over the remaining life of the straddle contract under consideration. It is assumed that the agent has \$100 to invest at the time of each trade. Noh, Engle, and Kane (1994) showed how to compute the rates of return for buying and selling straddles. When a straddle is sold, the agent is allowed to invest the proceeds in a risk-free asset. The rate of return from buying or selling straddles is computed as

$$RT_t = [100/(C_{t-1} + P_{t-1})] \cdot [C_t + P_t - C_{t-1} - P_{t-1}], \quad (19)$$

$$RT_t = [100/(C_{t-1} + P_{t-1})] \cdot [-(C_t + P_t - C_{t-1} - P_{t-1})] + 100r_f. \quad (20)$$

When a one-way transaction cost of \$0.05 per straddle (\$0.025 per option) is imposed, the net rate of return is computed as,

$$NRT_t = RT_t - [100/(C_{t-1} + P_{t-1})] \cdot 0.10 \quad (21)$$

The average daily rate of return from trading straddles is computed for each volatility forecast under all six test scenarios. The average daily profit streams are analyzed to determine if they are correlated with the market return and significantly different from zero. A close to zero correlation will mean that the trading profits do not possess any systematic risk and can therefore be considered as risk-free profits. The existence of risk-free profits would imply an inefficient options market.¹³

Empirical Results

The results in the "Comination Forecasts" section indicate that future volatility can be predicted, at least in a statistical sense. In this section, we explore the economic value of the forecasts and the implications for the efficiency of the Toronto 35 index options market.

Table 3 shows the average option prices computed by the different forecasting models. The implied volatility forecast leads to model prices of straddles higher than the actual market prices, which results in more buy deci-

Table 4
Straddle Test Results^a

Forecast	Closing prices No transaction costs			Closing prices \$0.05 transaction costs			Bid-ask prices ^c		
	Mean return (%)	<i>t</i> value	<i>R</i> ^{2b}	Mean return (%)	<i>t</i> value	<i>R</i> ²	Mean return (%)	<i>t</i> value	<i>R</i> ²
Ex post test results									
Historical volatility	2.1285	6.4693	0.0025	-0.0935	-0.2828	0.0025	-9.2363	-28.3742	0.0013
GARCH(1,1)	2.0791	6.3135	0.0039	-0.1429	-0.4317	0.0039	-9.3845	-28.4516	0.0028
EGARCH(1,1)	2.0692	6.2824	0.0070	-0.1528	-0.4613	0.0070	-9.3528	-28.2179	0.0040
Implied volatility	1.5521	4.6749	0.0021	-0.6699	-2.0097	0.0021	-9.9638	-31.4558	0.0016
Simple average	2.2147	6.7423	0.0037	-0.0073	-0.0220	0.0038	-9.4689	-29.0601	0.0024
Weighted average	2.3687	7.2331	0.0061	0.1467	0.4463	0.0062	-9.3262	-28.7124	0.0049
GARCH/IV	2.4509	7.4971	0.0040	0.2289	0.6970	0.0041	-9.1813	-27.3823	0.0060
Ex ante test results									
Historical volatility	0.2329	0.5658	0.0002	-1.9891	-4.8171	0.0001	-10.7119	-24.9205	0.0003
GARCH	0.5895	1.4334	0.0020	-1.6325	-3.9585	0.0001	-10.4532	-24.1978	0.0001
EGARCH	0.5197	1.2634	0.0001	-1.7023	-4.1224	0.0001	-10.5932	-24.3362	0.0001
Implied volatility	-0.0504	-0.1225	0.0002	-2.2724	-5.5113	0.0002	-10.5858	-25.5820	0.0004
Simple average	0.6941	1.6884	0.0001	-1.5279	-3.7052	0.0001	-10.3871	-24.1175	0.0001
Weighted average	0.5716	1.3899	0.0001	-1.6503	-3.9978	0.0001	-10.3216	-23.9786	0.0001
GARCH/IV	0.3697	0.8984	0.0005	-1.8523	-4.4897	0.0005	-10.9141	-24.6416	0.0003

^aThe straddle tests use 1,067 trading days from January 1991 to July 1995.

^bThe *R*² values provide a measure of the correlation between the straddle profits and the market's return (as measured by the return on the Toronto 35 index). A low *R*² indicates a very small correlation.

^cClosing prices are used to make the buy/sell decision each day. Once the buy/sell decision is made, the position is opened and closed at the bid-ask prices.

sions. All other forecasts, especially the GARCH/IV forecast, result in straddle prices lower than the actual straddle prices. Therefore, these forecasts lead to more sell decisions than buy decisions.

The top half of Table 4 shows the results of the ex post straddle trading tests. The straddle trading profits are expressed as the daily percentage rate of return per \$100 invested. The first three columns show the profits from trading under the assumption of no transaction costs. With the exception of the implied volatility forecast, all of the forecasts show significantly positive returns of more than 2% per day. The three combination forecasts earn the largest returns, followed by the GARCH forecasts, and then implied volatility.

Rates of return of 2% per day are impressive; *R*² values of less than 0.007 indicate that these returns are not correlated with the market's rate of return and that no risk adjustment is needed.¹⁴ Therefore, these are abnor-

mal returns. Even though they are only average returns and therefore are not certain, they do indicate that the volatility forecasts can indeed generate trading profits before transaction costs.

The middle three columns and the last three columns of the top panel of Table 4 show the results of the ex post tests when transaction costs are considered. Transacting at closing prices with an assumed cost of \$0.05 per straddle (incurred on both the opening and closing transactions) results in insignificantly negative returns for the GARCH and implied volatility forecasts. The weighted average combination forecast and GARCH/IV forecast earn insignificantly positive returns of 0.1467% and 0.2289% per day. Transacting at bid-ask option prices results in significantly negative returns for all forecasts. For example, the weighted average combination forecast and the GARCH/IV forecast now earn returns of -9.3262% and -9.1813% per day. The mag-

nitude of these negative returns provides some indirect evidence for the conjecture that transacting at bid and ask prices overstates the true size of the transaction spread.

The ex post tests indicate that our straddle trading rules are able to generate profits before transaction costs. To reach a definite conclusion about market efficiency, we now proceed to ex ante tests, which are presented in the bottom panel of Table 4. The first three columns show that the mean daily returns of the ex ante tests are considerably smaller than those of the ex post tests. For example, before transaction costs, the GARCH (1,1) forecast earns an average daily ex post return of 2.0791%; ex ante, the GARCH(1,1) forecast earns 0.5895% per day. This decrease reflects the fact that ex ante tests do not allow us to transact at the prices that are generating the profit signal. In other words, the market seems to be prompt in correcting price deviations.

Ex ante, none of the forecasts earn positive returns that are significant at the 5% level. The simple average combination forecast earns the largest daily return, 0.6941%, with a corresponding *t* statistic of 1.6884. The implied-volatility forecast earns a daily return of -0.0504%, which is lower than any of the other forecasts. The next six columns show the results when transaction costs are included in the analysis. The introduction of transaction costs results in significantly negative returns for all of the forecasts. Similar to the ex post tests, the negative returns are magnified when transactions are made at bid and ask prices. Clearly, we cannot earn abnormal returns when transaction costs are considered.

Although the volatility forecasts may be reasonably accurate from a statistical point of view and the ex post tests indicate our trading rule can generate some profits before transaction costs, the ex ante tests strongly suggest that abnormal returns cannot be earned. Therefore, we conclude that the Toronto 35 index options market is efficient with respect to the straddle trading rules employed in this study. Specifically, we conclude that the market is efficient with respect to GARCH conditional-variance forecasts, implied-volatility forecasts, and forecasts that combine the information contained in past stock index returns and option prices in the market.

Finally, it is interesting to note that the implied-volatility forecast earns the lowest profits, and is outperformed by all other forecasts, including the historical volatility, in almost all scenarios. This observation is consistent with the statistical evaluations of the forecasts presented at the end of the "Combination Forecasts" section. It is shown there that implied volatility is dominated by all other forecasts in terms of mean absolute error, mean absolute percent error, and root-mean-square error. The observation is also similar to that made by Canina

and Figlewski (1993) with the S & P 100 index options. Similar to Canina and Figlewski, we find that even the simple historical volatility is a better forecast than the implied volatility. There are many possible explanations. One is that investors are simply not rational and consistently make errors in forecasting future volatility; another is that the implied volatility is biased because of the nonlinear functional relation between the option price and the volatility. However, insofar as we use near-the-money options to imply volatilities, the bias, if any, should be small. (It is well known that in the Black-Scholes model, the option price is almost linear in the volatility when the exercise price is close to the underlying asset's price.)

It should also be pointed out that implied volatility is perhaps the least desirable variance measure to input into the Black-Scholes model as far as evaluating market efficiency is concerned. Because implied volatilities are derived from market prices of options and then substituted back to the same model to obtain theoretical prices, the model prices obtained will be, by design, close to the actual market prices. (This is very clear from Table 3.) As a result, it is doubtful that any trading strategy based solely on implied volatilities would lead to a fair assessment of market efficiency. For the reasons discussed above, studies employing only implied volatilities will tend to favour the hypothesis of the market being efficient. Since we have used a variety of other volatility forecasts, our study is immune to this criticism.

Conclusions

This study investigates the efficiency of the Toronto 35 index options market over the period from January 1991 to July 1995 by using a variety of volatility forecasting techniques. A delta neutral straddle trading strategy driven by different volatility forecasts is used to test the efficiency of this market. After transaction costs, no abnormal profits can be earned, regardless of whether the same day's or the following day's prices are used to initiate a trading position. In other words, both the ex post and the ex ante tests fail to detect abnormal profits. The results clearly indicate that the Toronto 35 index options market is efficient over the sample period considered. This is an important result; previous studies of Canadian stock options markets found strong evidence of inefficiency. It may be inferred that the options markets have caught up in terms of sophistication and become more efficient over the years. Given our findings, regulators and investors can be reassured that prices in the Toronto 35 index options market are fair and that it is not possible to earn systematic abnormal returns.

Implied volatility is widely viewed as one of the

best forecasts of volatility. Our results suggest otherwise. Investment professionals in Canada who use implied volatility to make investment decisions may want to consider alternative forecasting techniques such as GARCH or a combination of GARCH and implied volatilities.

References

- Akgriray, V. (1989). Conditional heteroskedasticity in time series of stock returns: Evidence and forecasts. *Journal of Business*, 62, 55-80.
- Berndt, E., Hall, B., Hall, R., & Hausman, J. (1974). Estimation and inference in nonlinear structural models. *Annals of Economic and Social Measurement*, 4, 653-665.
- Bhattacharya, M. (1983). Transaction data tests of the efficiency of the Chicago Board Options Exchange. *Journal of Financial Economics*, 12, 161-185.
- Black, F. (1976). Studies of stock price volatility changes. In *Proceedings of the Business and Economic Statistics Section*, (pp. 177-181). Washington, DC: American Statistical Association.
- Black, F., & Scholes, M. (1972). The valuation of option contracts and a test of market efficiency. *Journal of Finance*, 27, 399-417.
- Black, F., & Scholes, M. (1973). The pricing of options and corporate liabilities. *Journal of Political Economy*, 81, 737-654.
- Bollerslev, T. (1986). Generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics*, 31, 307-327.
- Bollerslev, T., Chou, R., & Kroner, K. (1992). ARCH modeling in finance. *Journal of Econometrics*, 52, 5-59.
- Brailsford, T., & Faff, R. (1996). An evaluation of volatility forecasting techniques. *Journal of Banking and Finance*, 20, 419-438.
- Canina, L., & Figlewski, S. (1993). The informational content of implied volatility. *Review of Financial Studies*, 6, 659-681.
- Chance, D. (1986). Empirical tests of the pricing of index call options. *Advances in Futures and Options Research*, 1, 141-166.
- Christie, A. (1982). The stochastic behavior of common stock variances: Value, leverage, and interest rate effects. *Journal of Financial Economics*, 10, 407-432.
- Chua, J., & Mooklelost, P. (1988). An empirical study of Trans-Canada options: A preliminary report. *Annual Meeting of the Administrative Sciences Association of Canada-Finance Division*, 1-12.
- Clemen, R. (1989). Combining forecasts: A review and annotated bibliography. *International Journal of Forecasting*, 5, 559-583.
- Corrado, C., & Miller, T. (1996). Efficient option-implied volatility estimators. *Journal of Futures Markets*, 16, 247-272.
- Day, T., & Lewis, C. (1988). The behavior of the volatility implicit in the prices of stock index options. *Journal of Financial Economics*, 22, 103-122.
- Day, T., & Lewis, C. (1992). Stock market volatility and information content of stock index options. *Journal of Econometrics*, 52, 267-287.
- Engle, R. (1982). Autoregressive conditional heteroskedasticity with estimates of the variance of United Kingdom inflation. *Econometrica*, 50, 987-1007.
- Galai, D. (1977). Tests of efficiency of the Chicago Board Options Exchange. *Journal of Business*, 50, 167-197.
- Galai, D. (1978). Empirical tests of boundary conditions for CBOE options. *Journal of Financial Economics*, 8, 179-201.
- Glosten, L., Jagannathan, R., & Runkle, D. (1993). On the relation between the expected value and the volatility of the nominal excess return on stocks. *Journal of Finance*, 48, 1779-1801.
- Halpern, P., & Turnbull, S. (1985). Empirical tests of boundary conditions for Toronto Stock Exchange Options. *Journal of Finance*, 40, 481-500.
- Harvey, C., & Whaley, R. (1992). Market volatility prediction and the efficiency of the S&P 100 index option market. *Journal of Financial Economics*, 31, 43-73.
- Heynen, R., Kemna, A., & Vorst, T. (1994). Analysis of the term structure of implied volatility. *Journal of Financial and Quantitative Analysis*, 29, 31-56.
- Lamoureux, C., & Lastrapes, W. (1993). Forecasting stock-return variance: Toward an understanding of stochastic implied volatilities. *Review of Financial Studies*, 6, 293-326.
- Mandron, A. (1988). Some empirical evidence about Canadian stock options. Part I: Valuation. *Canadian Journal of Administrative Sciences*, 5(2), 1-13.
- Nelson, D. (1991). Conditional heteroskedasticity in asset returns: A new approach. *Econometrica*, 59, 347-370.
- Noh, J., Engle R., & Kane, A. (1994). Forecasting volatility and option prices of the S&P 500 Index. *Journal of Derivatives*, fall, 17-30.
- Stein, J. (1989). Overreactions in the options market. *Journal of Finance*, 44, 1011-1023.
- Vasilellis, G., & Meade, N. (1996). Forecasting volatility for portfolio selection. *Journal of Business Finance and Accounting*, 23, 125-143.

Notes

1. During periods of high trading volume, contracts at strike price increments of 2½ points are sometimes introduced.
2. The extensive survey by Bollerslev, Chou, and Kroner (1992) demonstrates the popularity and usefulness of GARCH models.
3. The statistical software package RATS Version 4.2 is used to do the required estimations.
4. A GARCH-M model and a GARCH model with a conditional t density were also considered in this study. Because the estimated parameters and forecasts were not significantly different from those of the GARCH model, the results are not included in this paper.
5. The GJR-GARCH model by Glosten, Jagannathan, and

- Runkle (1993) was also considered in this study. Because the results were similar to those of the EGARCH model, they are not included in this paper.
6. Corrado and Miller (1996) state that implied volatilities derived from short-term, at-the-money options minimize any biases induced by stochastic volatility, price observation errors, and the bid-ask spread.
 7. Implied volatilities were derived from both closing option prices and bid-ask option prices. Because no significant difference was found between the two implied volatility series, only the volatility series implied in closing option prices is used in this study.
 8. Ideally, the new GARCH model should also be estimated on a daily rolling basis. However, using a rolling sample of 600 observations resulted in convergence problems. Therefore, the GARCH/IV model parameters are estimated over the full sample (1,743 observations) only.
 9. Treating the expectation as a constant may seem oversimplified; however, implied volatility is forward-looking in the sense that it represents the average volatility expected over the remaining life of the option under consideration. Treating the expectation as a constant also avoids the problem of choosing an appropriate process to model the implied volatility.
 10. The average call option price is \$2.65; the average put option price is \$2.35. For prices in this range, the appropriate cost listed on the *TD Greenline Investment Fee Schedule* is \$0.025 per option.
 11. A long (short) straddle is formed by buying (selling) a call option and a put option that have the same strike price and expiration date.
 12. Among the 1,067 trading days used in the straddle tests, there are only about 50 days when a transaction has to be closed in 2 days. For all the remaining days, no delay in completing a transaction is experienced. It should be pointed out that the prices for the call and put in a straddle are not likely to be observed simultaneously. As a result, the accuracy of trading profits for a particular straddle may not be high. However, insofar as the profits reported are averaged over the sample period, the bias caused by the liquidity difficulty is hopefully canceled out. Finally, it should be kept in mind that our trading strategy can only shed light on the market efficiency over a daily sampling interval. Whether trading profits exist within a trading day can only be resolved by examining intra-day data.
 13. It should be noted that our market-efficiency tests are actually tests of the joint hypothesis that the option-pricing model is correct and that the market is efficient. Given the strong empirical support for the Black-Scholes model in pricing short-term, at-the-money options, we can be fairly confident that any biases introduced by an incorrect model specification will be minimal.
 14. The daily rate of return on the Toronto 35 index is used as a proxy for the market's rate of return.