



Predicting Volatility in the Foreign Exchange Market

Philippe Jorion

The Journal of Finance, Volume 50, Issue 2 (Jun., 1995), 507-528.

Stable URL:

<http://links.jstor.org/sici?sici=0022-1082%28199506%2950%3A2%3C507%3APVITFE%3E2.0.CO%3B2-K>

Your use of the JSTOR archive indicates your acceptance of JSTOR's Terms and Conditions of Use, available at <http://www.jstor.org/about/terms.html>. JSTOR's Terms and Conditions of Use provides, in part, that unless you have obtained prior permission, you may not download an entire issue of a journal or multiple copies of articles, and you may use content in the JSTOR archive only for your personal, non-commercial use.

Each copy of any part of a JSTOR transmission must contain the same copyright notice that appears on the screen or printed page of such transmission.

The Journal of Finance is published by American Finance Association. Please contact the publisher for further permissions regarding the use of this work. Publisher contact information may be obtained at <http://www.jstor.org/journals/afina.html>.

The Journal of Finance

©1995 American Finance Association

JSTOR and the JSTOR logo are trademarks of JSTOR, and are Registered in the U.S. Patent and Trademark Office. For more information on JSTOR contact jstor-info@umich.edu.

©2002 JSTOR

Predicting Volatility in the Foreign Exchange Market

PHILIPPE JORION*

ABSTRACT

Measures of volatility implied in option prices are widely believed to be the best available volatility forecasts. In this article, we examine the information content and predictive power of implied standard deviations (ISDs) derived from Chicago Mercantile Exchange options on foreign currency futures. The article finds that statistical time-series models, even when given the advantage of "ex post" parameter estimates, are outperformed by ISDs. ISDs, however, also appear to be biased volatility forecasts. Using simulations to investigate the robustness of these results, the article finds that measurement errors and statistical problems can substantially distort inferences. Even accounting for these, however, ISDs appear to be too variable relative to future volatility.

IT IS WIDELY BELIEVED that the volatility implied in option prices is the market's best estimate of future volatility. After all, if it were not, one could devise a trading strategy that could generate profits by identifying mispriced options.

The purpose of this article is to investigate the information content and predictive power of volatility implied in options on foreign currencies. The paper focuses on options on currency futures, traded on the Chicago Mercantile Exchange (CME). Because of the depth and liquidity of CME futures and options, traded side-by-side in the same market, implied standard deviations (ISDs) allow for clean tests of the predictive power of implied volatilities.

Early studies of the information content of ISDs have generally found that these contain substantial information for future volatility. Latané and Rendleman (1976), Chiras and Manaster (1978), and Beckers (1981), for example, regress future volatility on the weighted implied volatility across a broad sample of Chicago Board Options Exchange (CBOE) stocks, and find that options contain volatility forecasts that are more accurate than historical measures.¹ These studies, performed shortly after the 1973 beginning of the CBOE option market, could only use a relatively short time span, and therefore focused on cross-sections rather than time-series predictions.

* Jorion is from the Graduate School of Management, University of California at Irvine. Thanks are due to David Bates, Hendrik Bessembinder, Michael Brennan, Stephen Figlewski, Steven Heston, two referees, and seminar participants at UC-Irvine, UCLA, Georgetown University, the Université Libre de Bruxelles, the University of Maryland, the University of Wisconsin-Madison, and the Cornell conference on derivatives for useful comments. I am also grateful to the Institute for Quantitative Research in Finance for financial support.

¹ The slope coefficients, however, were generally around 0.5, instead of unity.

More recently, research has turned to the analysis of volatility in a time-series framework. Scott and Tucker (1989) report some predictive ability in ISDs measured from PHLX currency options, but their methodology does not allow formal tests of hypotheses.² Day and Lewis (1992) analyze options on the S&P 100 Index from 1983 to 1989, and find that the ISD has significant information content for weekly volatility, although not necessarily higher than that of time-series models. This approach, however, ignores the term structure of volatility since the return horizon is not matched with the life of the option. To address this problem, Canina and Figlewski (1993) regress the volatility over the remaining contract life against the implied volatility of S&P 100 Index options over 1983 to 1986. They report that ISDs have little predictive power for future volatility, and are significantly biased forecasts. Furthermore, option volatilities appear to be even worse than simple historical measures. Finally, Lamoureux and Lastrapes (1993) focus on individual stock options, carefully measuring prices and matching the forecast horizon, and find that historical time-series contain predictive information over and above that of implied volatilities. They view their result as a rejection of the joint hypotheses of market efficiency and of the Black-Scholes (BS) class of option pricing models.

As with all efficiency tests, these results can be interpreted either as indicative of misleading test procedures, or as indicative of inefficient processing of information by option markets. The present article investigates whether these results carry over to the currency option markets, and provides a detailed analysis of the impact of faulty test procedures due to measurement errors or inappropriate statistical inferences.

Previous work on the information content of ISD has paid little attention to the effect of measurement errors in reported prices. These can cause major problems with estimation of volatility. S&P 100 options, for instance, although actively traded, suffer from several shortcomings. Even with time-stamped quotes on the underlying S&P 100 Index, it is unlikely that all 100 underlying prices reflect trades that are simultaneous with the option trade. Given the sheer size of the underlying portfolio, some of the quotes must be "stale," which creates measurement problems for the option volatility. In addition, the arbitrage between options and the underlying asset is difficult to implement because of the transaction costs involved in going long or short 100 stocks simultaneously. Because of the lower transaction costs involved when trading options, news tends to be more rapidly disseminated in option prices than in the prices of underlying stocks, which causes measurement errors in the volatility. Finally, as in all markets, quoted prices may be at the bid, the ask, or in between, which creates an additional errors-in-variables problem. Harvey and Whaley (1991) also find that differences in closing

² Scott and Tucker (1989) present one OLS regression with 5 currencies, 3 maturities, and 13 different dates. Because of correlations across observations, the usual OLS standard errors are severely biased, thereby invalidating hypothesis tests.

times, bid-ask spreads, and to some extent infrequent trading contribute to biases in the daily autocorrelation of ISDs.

Another potential of source of biases is an inappropriate option pricing model. In their original derivation, Black and Scholes (1973) assume a nonstochastic volatility for the underlying asset. Apparently, there is some inconsistency in recovering an implied volatility from inverting the BS model, and then proceeding to study the stochastic behavior of volatility. In theory, one would want to invert an option pricing model consistent with stochastic volatility. However, no previously published article has done this, for three reasons. First, current stochastic volatility models involve costly numerical simulations. It is necessary, not only to compute the option price as a function of all parameters, but also iterate to find an implied instantaneous volatility. Second, these models assume a specific time-series process for the volatility. This model may be misspecified and involves the estimation of additional parameters, which introduces a supplementary source of error. Finally, for the short-term at-the-money options considered here, the BS model is very close to linear in the average volatility, and yields estimates virtually identical to those of a stochastic volatility model.³

The present study focuses on options on foreign currency futures. These are actively traded contracts, with average notional volume now approaching \$5 billion daily. In addition, both the underlying asset and the options are traded side-by-side on the CME, and close at the same time. Because of the low cost of transacting between the two markets, options and futures prices are likely to be highly synchronized, which alleviates the problem of nonsynchronous quotes afflicting S&P 100 Stock Index options. As an added advantage, all CME closing quotes are carefully scrutinized by the exchange because they are used for daily settlement, and therefore less likely to suffer from clerical measurement errors. The currencies selected for this study are the three most actively traded currencies: the German deutsche mark (DM), the Japanese yen (JY), and the Swiss franc (SF).

This article investigates both the informational content and the predictive power of volatility implied in option prices. Informational content is measured in terms of the ability of the explanatory variable to forecast 1-day volatility. Tests of predictive power, in contrast, focus on the volatility over the remaining days of the option contract. ISDs are pitted against time-series models such as a moving average and Generalized Autoregressive Conditional Heteroskedastic model (GARCH). In order to gain maximum efficiency within a limited sample period, daily observations are used, and standard errors are corrected for the data overlap.

Even in the highly synchronized currency futures and options markets, ISDs are still affected by measurement errors such as bid-ask spreads. In addition, the predictive power regressions suffer from small sample biases,

³ Fleming (1993) summarizes the theoretical argument for using the Black-Scholes ISD as a rational forecast of the average volatility over the life of the option. Section II below also shows that the biases are small for such options.

which render the interpretation difficult. The extent of these problems is illustrated by Monte Carlo simulations that show that inferences can be affected by nonnegligible biases. The simulations also provide a method to gauge whether ISDs can be considered the best available forecasts of future volatility.

This article is organized as follows. The basic regression setup is presented in Section I. Section II shows how the two option pricing models are used to compute the implied volatility. Section III describes the data and basic statistics. Predictive regressions are presented in Section IV and interpreted in light of the Monte-Carlo simulations in Section V. Finally, the last section contains some concluding observations.

I. Testing Predictive Power

The ISD is widely believed to be the best available forecast of the volatility of returns over the remaining contract life. Define $\sigma_{t,T}$ to be the realized volatility over the remaining life of the contract, measured from day t to day T . In what follows, the future daily variance is taken from the arithmetic average of squared returns over future trading days, $\sigma_{t,T}^2 = (1/(T-t))\sum_{i=1}^{T-t} R_{t+i}^2$, without adjustment for the mean.⁴

The predictive power of a volatility forecast can be estimated by regressing the realized volatility on forecast volatility:

$$\sigma_{t,T} = a + b\hat{\sigma}_t + \varepsilon_{t,T} \quad (1)$$

where $\hat{\sigma}_t = \sigma_t^{ISD}$ is the volatility forecast measured on day t , taken as the implied volatility. As in typical tests of market efficiency, if the ISD were the best available forecast of future volatilities, we would expect the intercept to be zero and the slope coefficient to be unity.

This framework can be extended to comparisons of the predictive power of the ISD with that of a simple time-series model, σ_t^{TS} . The latter can be taken as a simple moving average, estimated for example over a 20-day window, $\sigma_t^2 = (1/20)\sum_{i=1}^{20} R_{t-i+1}^2$, or from a more sophisticated GARCH model. From the voluminous literature documenting time-variation in second moments in the foreign exchange market, the paper adopts the GARCH(1, 1) specification, because it is a parsimonious representation that seems to fit the data relatively well.⁵ The GARCH model, developed by Engle (1982) and extended by Bollerslev (1986), posits that the variance of returns follows a predictable process, driven by the latest squared innovation and by the previous conditional variance:

$$\tilde{R}_t = \mu + \tilde{r}_t, \tilde{r}_t \sim N(0, h_t), h_t = \alpha_0 + \alpha_1 r_{t-1}^2 + \beta h_{t-1} \quad (2)$$

⁴ With daily data, the average term $E(R^2)$ dominates the term $E(R)^2$ by a typical factor of 700 to one. Therefore, ignoring expected returns is unlikely to cause a perceptible bias in the volatility estimate.

⁵ For applications, see for instance Hsieh (1989), and Giovannini and Jorion (1989).

where R_t is the nominal return, \tilde{r}_t is the de-meaned return and h_t its conditional variance, measured at time t . To insure invertibility, the sum of parameters $(\alpha_1 + \beta)$ must be less than unity; when this is the case, the unconditional, long-run, variance is given by $\alpha_0/(1 - \alpha_1 - \beta)$.

When pitted against a time-series model, as in the following regression,

$$\sigma_{t,T} = a + b_1 \sigma_t^{ISD} + b_2 \sigma_t^{TS} + \varepsilon_{t,T}, \quad (3)$$

we would expect σ^{TS} to have no incremental predictive power beyond that in σ^{ISD} . In other words, the coefficient b_2 should be close to zero.

The previous specification precisely matches the ISD with the volatility over the remaining contract life, which is appropriate if one wants to test the forecasting accuracy of implied volatilities. Another, related, issue is the *information content* of the daily ISD for volatility over the next day, which can be tested as before,

$$\sqrt{R_{t+1}^2} = a + b \sigma_t^{ISD} + \varepsilon_{t+1}. \quad (4)$$

This approach is similar in spirit to the work of Day and Lewis (1992), who analyze the information content of the ISD with respect to future *weekly* stock index volatility. This setup, however, addresses the question of whether there is *some* useful information in σ^{ISD} , rather than whether σ^{ISD} is the best available forecast of future volatility. Since the forecast horizon is not matched with the realized returns, we only require the slope coefficient b to be positive, and not necessarily unity.

Because currency options are relatively recent financial instruments, the limited data span makes it important to fully utilize all the information in the sample. Therefore, daily data are used in the regressions. With horizons up to three months, however, using all the daily data introduces overlaps in the error terms, which causes downward bias in the usual ordinary least squares (OLS) standard errors. The point estimates of the coefficients, however, are still consistently estimated. Intuitively, the reason for this bias is that OLS estimation assumes that each day brings a completely independent observation, and therefore that the large number of observations permits precise estimation of the coefficients, when in fact most of the new observations are redundant.

Hansen (1982) provides a correction that extends White's (1980) method to deal with heteroskedasticity and properly deals with serial correlation.⁶ The corrected Hansen-White (HW) variance-covariance matrix of estimated coefficients is given by

$$\Sigma = (X'X)^{-1} \Omega (X'X)^{-1}, \quad (5)$$

⁶ Hansen and Hodrick (1980) first apply this method to test the predictive power of three-month forward rates measured at monthly intervals. Canina and Figlewski (1993) apply this method to volatility forecasts.

where $\Omega = E[X' \varepsilon \varepsilon' X/T]$ is consistently estimated, using the OLS residuals $\hat{\varepsilon}$, by

$$\hat{\Omega} = \sum_t \hat{\varepsilon}_t^2 X_t' X_t + \sum_s \sum_t Q(s, t) \hat{\varepsilon}_s \hat{\varepsilon}_t (X_s' X_s + X_t' X_t), \quad (6)$$

with $Q(s, t)$ defined as an indicator function equal to unity if there is overlap between returns at s and t , and zero otherwise. Note that in the case where the residuals are homoskedastic and do not overlap, $E[\varepsilon_t^2] = s^2$, and $Q(s, t)$ is always zero, so that the covariance matrix collapses to the usual OLS covariance matrix $s^2(X'X)^{-1}$.

II. Computing Implied Volatilities

Implied volatilities are derived from the Black (1976) model for European options on futures,

$$c = [FN(d_1) - KN(d_2)]e^{-r\tau}, \quad d_1 = \frac{\ln(F/K)}{\sigma\sqrt{\tau}} + \frac{\sigma\sqrt{\tau}}{2}, \quad d_2 = d_1 - \sigma\sqrt{\tau}, \quad (7)$$

where F is the futures rate, K is the strike price, τ is the time to option expiration, r is the risk-free rate, taken as the Eurodollar rate, and σ the volatility. For a given option price, inverting the pricing model yields an *implied* standard deviation. Because Beckers (1981) shows that using only at-the-money options was preferable to various other weighting schemes, only at-the-money calls and puts are considered here. In addition, these are the most actively traded, and therefore least likely to suffer from nonsimultaneity problems. On any given day, the ISD is computed as the arithmetic average of that obtained from the two closest at-the-money call and put options. Averaging a call and a put alleviates some of the measurement problems.

Since CME options are of the American type, using a European model introduces a small upward bias in the estimated volatility. This bias is generally considered small for short-maturity options.⁷ For instance, with typical parameter values,⁸ using a European model overestimates a 12

⁷ As shown by Whaley (1986), the bias depends on the level of the interest rate, the volatility, the time to expiration, and the degree to which the option is in-the-money. Shastri and Tandon (1986) use numerical procedures to show that biases in measured implied volatilities are generally minor for short-term at-the-money option. However, some error is also introduced because the nearest at-the-money option changes from day to day, which creates some spurious movements in ISDs.

⁸ With a futures price of 50 cents, a strike price of 50, a U.S. interest rate of 6 percent, 50 calendar days to expiration, and a true volatility of 12 percent, the value of an American and European call is 0.8799 and 0.8786, respectively. Inverting the American call value using a European model yields an apparent volatility of 12.02 percent. With the same parameters but 95 days to expiration, the estimated volatility is 12.04 percent. With 5 days to expiration, it is 12.00 percent.

percent true volatility by reporting a value of about 12.02 percent. The difference, however, is much less than typical bid-ask spreads, when quoted in terms of volatility, and in any event tends to bias the slope coefficient *downward* by a very small amount. Given the direction of the bias, findings of significance indicate that rejection of the hypothesis of no forecasting ability for ISDs would also occur with volatilities derived from a more accurate American model.⁹

Another potential misspecification is that the BS model is, *stricto sensu*, inconsistent with stochastic volatilities. If volatility changes in a deterministic fashion, the implied volatility can be construed as an average volatility over the remaining life of the option. But if volatility is stochastic, the arbitrage argument behind the BS option pricing model fails.

Recent papers by Hull and White (1987), Scott (1987), and Wiggins (1987) have examined the pricing of options on assets with stochastic volatility. The general approach to pricing options in these papers is to treat the volatility as a random state variable. In order to derive tractable results, the innovations in volatility and returns are generally assumed to be uncorrelated; prices are then calculated by Monte Carlo simulation. Scott (1988) and Chesney and Scott (1989), for instance, present a careful empirical analysis of the random variance model (implemented on a Cray supercomputer), and find that the random variance model actually provides a worse fit to market prices than the Black-Scholes model using ISDs. For U.S. stock options, differences are only on the order of 2 cents, which is much lower than typical bid-ask spreads of 5 to 25 cents. Duan (1995) extends the risk neutral valuation to the case where logarithmic returns follow a GARCH process. Under some combination of preferences and distribution assumptions, he derives a GARCH option pricing model, but the magnitude of the bias, computed by simulations, is very small, at most 10 to 15 cents for an option on a \$100 underlying asset.

Because of the computational costs involved in pricing options, no published research has ever recovered the implied (instantaneous) standard deviation from a stochastic volatility model.¹⁰ It is only recently that Heston (1993) has developed a closed-form solution that efficiently computes option values under stochastic volatility. These models, however, introduce additional difficulties. The time-series process for the volatility may be misspecified. Also, the models involve the estimation of additional parameters, which introduces a supplementary source of error. Finally, the gain from these stochastic volatility models is limited because the mispricing is very small for

⁹ With the above numbers, using a European model decreases the estimated slope coefficient from unity to 0.998. This bias, however, cannot account for the slope coefficients below.

¹⁰ Melino and Turnbull (1990), for instance, compare option prices derived from Black-Scholes and a stochastic volatility model, using parameters derived from the time-series process, and find that the stochastic volatility model provides a better fit to options than the standard model using historical volatility. They do not, however, consider a Black-Scholes model with implied volatility.

short-term at-the-money options.¹¹ Thus, the BS approach provides a good working approximation to more complicated stochastic volatility models for the options considered in this paper.

III. Data and Descriptive Statistics

The data are taken from the CME's closing quotes for currency futures and options on futures. The currencies covered are the German deutsche mark (DM), Japanese yen (JY), and Swiss franc (SF). These are the most active contracts on the CME.¹² For the DM, the sample period covers January 1985 to February 1992, which represents more than 7 years of daily data, or about 1810 observations. The data start in July 1986 for the JY, and in March 1985 for the SF.

We now turn to the timing of observations and the issue of matching ISDs with future volatility. For each option contract, the implied volatility is matched with the sequence of price movements on the underlying futures contract until option expiration. The "realized"—or "future"—volatility is then taken from the variance of continuously compounded futures returns.¹³ In what follows, the observation interval is taken as a trading day. If volatility was primarily the result of exogenous information, and if this flow was constant for every day, the weekend variance should be three times the weekday variance. It is well-known, however, that the variance over weekends is only slightly higher than over weekdays. Baillie and Bollerslev (1989), for instance, report that the average daily variance is 0.448 for the DM/\$ rate, and that the Friday-to-Monday variance is 0.569, which is closer to the daily variance than to the variance over three weekdays. As suggested by French (1984), consistent with these findings, the volatility over the remaining life of the contract was converted to a daily volatility using the number of trading days, instead of calendar days. Annualized volatility measures are obtained by multiplying by the square root of 252, which is the number of trading days in a calendar year. Interest rates, as always, are associated with calendar days.

Contracts expirations follow the usual March-June-September-December cycle. On the first day of the expiration month, which is the time around which most rollovers into the next contract occur, the option series switches

¹¹ An earlier version of the paper provides tests using Heston's stochastic volatility model, and finds that the model cannot explain the bias in implied volatility reported below. A major difficulty in the implementation of this model is that it requires knowledge of additional parameters for the volatility process. Some of these can be derived from the time-series of exchange rates, using the GMM method proposed by Bates and Penacchi (1991), but are subject to estimation error.

¹² Over the period 1985 to 1991, the average daily volume of futures contracts was 30,340 for the deutsche mark, 22,290 for the Japanese yen, 21,850 for the Swiss franc, 11,540 for the British pound, and 4,150 for the Canadian dollar.

¹³ The horizon of option contracts varies from 3 to 100 calendar days, or from 3 to 70 trading days.

into the next quarterly contract. This generates a time-series of continuous 1-day returns, implied volatility, and matching realized volatility.

Table I provides a brief description of the data. Means, standard deviations, and autocorrelations are presented for returns and different measures of volatility: next day volatility, volatility realized over the contract life, option implied volatility, and the two time-series models.

As expected, daily returns appear to be uncorrelated; for the DM, their standard deviation is 0.787 percent per day, or 12.5 percent per annum. The 1-day volatility, in contrast, shows small positive but significant autocorrelation at various lags, which suggests that volatility is persistent.

Looking next at statistics for the realized (future) volatility, we find numbers that are typical of an averaged series with data overlap: very high autocorrelation and smaller standard deviation, about 0.197, than daily volatility, which is about 0.534 for the DM. There is, however, still some positive autocorrelation beyond the 100-day lag, which is the longest possible data overlap.

Finally, the implied volatility displays a similar pattern of autocorrelation; the lower standard deviation of 0.142 for the DM is consistent with the fact that the latter series is an expectation of the previous series. The two series have very similar average values: for instance, 0.773 and 0.745, respectively, for the DM. Based on this table, there is no indication that ISDs systematically underestimate or overestimate future volatility.

The results are qualitatively similar across currencies. It should also be kept in mind, when comparing regressions across these three currencies, that the results are not independent since the three exchange rates are positively correlated.¹⁴

Figure 1 displays the time-variation in σ^{ISD} for the DM, measured in percent per annum, as well as the corresponding realized volatility. The volatility averages about 12 percent per annum and displays substantial time-variation, which seems to be reflected in both series.

Estimates of the GARCH(1, 1) process are presented in Table II. In line with previous research, the GARCH model is highly significant, with a $\chi^2(2)$ statistic exceeding 60 for all currencies. This is much higher than the 1 percent upper fractile of the chi-square, which is 9.2. There is no question, therefore, that realized volatility does change over time. The process is persistent but also stationary, with values of $(\alpha_1 + \beta)$ around 0.96 for the DM and SF. This number implies that a shock to the variance has a half-life of $\ln(0.5)/\ln(0.96)$, which is about 17 days.

IV. Forecasting Regressions

We first test the information content of the ISD for next day volatility with regression equation (4). Results are presented in Table III, with one panel for

¹⁴ The pairwise correlations in the futures price change over the 1985-1992 period are: DM/JY, 0.705, DM/SF, 0.898, and JY/SF, 0.682. The DM and SF thus move more closely to each other.

Table I
Descriptive Statistics: Returns and Volatility Measures

Means, standard deviations, and autocorrelation coefficients for 1-day returns on currency futures expressed in percent, 1-day volatility, future volatility (over remaining life of associated option contract), Implied Standard Deviation (ISD) from option prices, MA(20) volatility (from a moving average over the past 20 days), GARCH conditional volatility for the next day. Currencies are the German deutsche mark (DM), Japanese yen (JY), and Swiss franc (SF). Periods end on February 28, 1992, and start in January 1985 (DM), July 1986 (JY), and March 1985 (SF).

	Mean	Std. Dev.	Autocorrelation									
			1	2	3	4	5	10	20	100	252	
DM												
1-Day return	0.030	0.787	0.012	-0.025	0.012	-0.033	0.015	-0.007	0.036	-0.023	-0.024	
1-Day volatility	0.578	0.534	0.063	0.045	0.067	0.079	0.074	0.132	0.051	0.045	0.029	
Future volatility	0.745	0.197	0.969	0.946	0.919	0.898	0.876	0.769	0.583	0.122	0.057	
ISD volatility	0.773	0.142	0.973	0.950	0.931	0.913	0.898	0.824	0.675	0.230	0.240	
MA(20) volatility	0.750	0.240	0.977	0.953	0.927	0.900	0.871	0.708	0.306	0.162	0.158	
GARCH volatility	0.778	0.156	0.948	0.900	0.855	0.815	0.777	0.611	0.334	0.161	0.153	
JY												
1-Day return	0.008	0.703	0.005	-0.033	-0.002	-0.021	0.028	0.048	-0.004	-0.006	0.012	
1-Day volatility	0.509	0.484	0.117	0.062	0.050	0.032	0.117	0.021	0.017	-0.031	0.016	
Future volatility	0.653	0.186	0.967	0.937	0.906	0.880	0.852	0.719	0.512	-0.106	-0.241	
ISD volatility	0.658	0.116	0.959	0.922	0.887	0.855	0.829	0.702	0.461	-0.162	-0.188	
MA(20) volatility	0.674	0.208	0.971	0.937	0.901	0.863	0.828	0.601	0.158	-0.102	-0.060	
GARCH volatility	0.705	0.125	0.828	0.689	0.577	0.491	0.440	0.204	0.063	-0.046	0.005	
SF												
1-Day return	0.029	0.860	0.005	-0.022	0.010	-0.031	0.019	0.006	0.027	-0.012	0.001	
1-Day volatility	0.645	0.567	0.015	0.006	0.048	0.061	0.073	0.099	0.042	0.022	0.034	
Future volatility	0.809	0.182	0.964	0.934	0.905	0.881	0.854	0.729	0.510	0.055	0.042	
ISD volatility	0.803	0.134	0.972	0.946	0.923	0.905	0.888	0.805	0.649	0.133	0.229	
MA(20) volatility	0.830	0.233	0.975	0.949	0.924	0.897	0.870	0.704	0.300	0.088	0.162	
GARCH volatility	0.854	0.126	0.950	0.906	0.868	0.832	0.800	0.643	0.366	0.091	0.157	

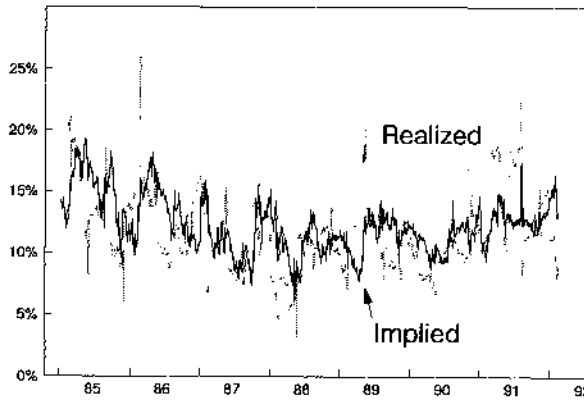


Figure 1. Implied and realized volatility: Mark. Time-series of volatility implied from option prices and of volatility subsequently realized over the life of the option contract.

each currency. Since the dependent variables do not overlap, standard errors are simply computed by OLS. The table shows that the ISD contains a substantial amount of information for currency movements the following day. In the first panel, the slope coefficient is 0.852 for the DM, and highly significant. Similarly high coefficients are obtained for the JY and SF. Dividing the sample period into two subperiods leads to similar conclusions.

Forecasting with a time-series model also produces significant results. The slope of the 20-day moving average is 0.395, and that of the GARCH 1-day forecast is 0.598 for the DM; both are significant for all currencies. The GARCH model appears less biased than the simple moving average model, although its explanatory power, expressed in terms of R^2 , is slightly lower. This is in spite of the fact that the parameters of the GARCH models were estimated over the whole period 1985 to 1992, and that we should consequently expect this model to outperform the MA(20) volatility. Although the GARCH parameters are generally stable over time, there is no guarantee that market participants would have been able to use the GARCH model as early as in 1985 (especially since the model was not yet published).

The next lines in the DM panel pit the ISD against the time-series models. There is no information content of the time-series models beyond that in options for all three currencies. These results are reassuring because they indicate that option traders form better expectations of risk over the next day than statistical models, even when the latter are based on ex post parameter values. To some extent, these results were expected, since time-series models are unable to account for events such as regular announcements of macroeconomics indicators, meeting of G-7 finance ministers, and so on. Because the timing of these events is known by the foreign exchange market, we would expect options to provide better forecasts than simple time-series models.

Predictability, however, implies that the ISD is an unbiased forecast over the remaining option life, not over the next day. There could very well exist a

Table II
Estimation of GARCH(1, 1) Process

$$\tilde{R}_t = \mu + \tilde{r}_t, \tilde{r}_t \sim N(0, h_t)$$

$$h_t = \alpha_0 + \alpha_1 r_{t-1}^2 + \beta h_{t-1}$$

where R_t is defined as the return on currency futures, and h_t is the conditional variance of the innovations. Currencies are deutsche mark (DM), Japanese yen (JY), and Swiss franc (SF). Periods end on February 28, 1992, and start in January 1985 (DM), July 1986 (JY), March 1985 (SF). Asymptotic t -statistics in parentheses. The χ^2 statistic tests the hypothesis of significance of added GARCH process. Returns are measured in percent per day.

Currency	Model	μ	α_0	α_1	β	Long-Run Volatility (% pa)	Log-Likelihood	$\chi^2(2)$ (p -value)
DM	Normal	0.0302 (1.63)	0.6190 (30.07)			12.49	6198.05	
	GARCH	0.0287 (1.65)	0.0268 (4.50)	0.0757 (5.87)	0.8828 (79.37)	12.77	6253.04	109.98
JY	Normal	0.0082 (0.44)	0.4940 (26.72)			11.16	5057.14	
	GARCH	0.0088 (0.49)	0.0586 (7.72)	0.0965 (5.08)	0.7870 (46.76)	11.26	5088.01	61.73
SF	Normal	0.0293 (1.43)	0.7400 (29.76)			13.66	5913.11	
	GARCH	0.0224 (1.14)	0.0322 (4.64)	0.0521 (4.38)	0.9042 (88.32)	13.32	5944.40	62.58

Table III
Information Content Regressions

$$\sqrt{R_{t+1}^2} = a + b\hat{\sigma}_t + \varepsilon_{t+1}$$

where R_{t+1} , the futures return over the next day, is regressed against the volatility forecast $\hat{\sigma}_t$. This includes the implied standard deviation (ISD) from option prices, a moving average (MA) with a moving window of 20 trading days, and the GARCH time-series model, which is the conditional volatility for the next day based on parameters in Table II. Periods end on February 28, 1992, and start in January 1985 (DM), July 1986 (JY), and March 1985 (SF). OLS standard errors are in parentheses.

Currency	a	Slopes On			R ²
		ISD	MA(20)	GARCH	
DM	-0.080 (0.068)	0.852* (0.086)			0.0515
	0.282* (0.041)		0.395* (0.052)		0.0315
	0.113 (0.063)			0.598* (0.079)	0.0304
	-0.069 (0.069)	0.784* (0.126)	0.055 (0.074)		0.0518
	-0.095 (0.070)	0.785* (0.123)		0.085 (0.112)	0.0518
	-0.006 (0.072)	0.783* (0.108)			0.0355
	0.283* (0.043)		0.335* (0.061)		0.0206
JY	0.119 (0.075)			0.560* (0.107)	0.0190
	0.001 (0.072)	0.678* (0.140)	0.093 (0.079)		0.0365
	-0.076 (0.083)	0.668* (0.128)		0.210 (0.125)	0.0374
	-0.041 (0.080)	0.854* (0.099)			0.0406
	0.335* (0.049)		0.373* (0.057)		0.0235
SF	0.071 (0.091)			0.672* (0.106)	0.0223
	-0.039 (0.083)	0.840* (0.150)	0.011 (0.086)		0.0406
	-0.040 (0.092)	0.858* (0.148)		-0.005 (0.157)	0.0406

* Significantly different from zero at the 5 percent level.

“term structure” of volatility if the market knows that information will be released at specific points in the future. With monthly trade figure announcements, for instance, the volatility could be higher on the third Thursday of every month, and lower in surrounding days. Because the volatility is not

necessarily constant over all contract days, it may not be appropriate to treat the current ISD as a forecast of next day's volatility.

The predictive power hypothesis is tested using equation (1), with results presented in Table IV. Here, the GARCH forecast is obtained by successively solving for the expected variance for each remaining day, then averaging over all days. Comparing to Table III, we observe that the R^2 s are systematically higher than before, which was to be expected, since there is more noise in daily volatility than in an average measure.

Table IV shows that ISDs contain a substantial amount of information for future volatility. The slope coefficient is 0.547 for the DM, with an HW t -statistic of 3.96 for the hypothesis that $\beta = 0$, which is significantly higher than zero. The results are remarkably similar across currencies. ISDs, however, also appear to be biased predictors of future volatility. The estimate of the slope coefficient is less than unity, and the associated HW t -statistic is 3.29 for the hypothesis that $1 - \beta = 0$, which is statistically significant. The slope coefficient less than unity combined with a positive intercept suggests that ISDs are too volatile: when ISDs are high relative to the average, they should be scaled down; when ISDs are low relative to the average, they should be brought up toward the intercept.

For all currencies, the MA(20) and GARCH time-series forecasts have lower explanatory power than ISDs. As in the case of the information content regressions in Table III, the slope coefficients become smaller and insignificant when using both the ISD and a time-series model in the same regression. The R^2 of the regressions is also noticeably higher when including ISDs. For the DM, for example, the R^2 increases from 0.0499 for the GARCH model to 0.1599 when including the ISD. This is a substantial improvement. Results for the SF and the JY are very similar. In each case, the additional coefficient on the MA(20) or the GARCH volatility is not significant; most of the information content lies in ISDs. These results indicate that options provide informative forecasts of future volatility that are superior to those of time-series models.

These results are in sharp contrast to those of Canina and Figlewski (1993), who report slope coefficients on ISDs and on a MA(60) model of 0.229 and 0.573, respectively. Combining both forecasts in the same regression, they find that the ISD coefficient drops to 0.077 and becomes insignificant. The poor performance of option volatilities, therefore, is sharply reversed in the foreign currency option market.

Since it is difficult to argue that option traders are smarter in the foreign exchange market than in the U.S. stock market, the most logical interpretation of the differences in the results is that S&P 100 Index option ISDs have been measured with substantial error because of stale prices and because of the difficulty of arbitraging between the option and underlying stock markets.

As an illustration of the seriousness of the problem, consider the impact of changes in underlying asset prices on the implied volatility. For a given

Table IV
Predictability Regressions

$$\sigma_{t,T} = \alpha + b\hat{\sigma}_t^T + \varepsilon_{t,T}$$

where $\sigma_{t,T}$, the volatility over the remaining life of the option contract, is regressed against the volatility forecast $\hat{\sigma}_t^T$. This includes the implied standard deviation (ISD) from option prices, a moving average (MA) with a moving window of 20 trading days, and the GARCH time-series model, which is the conditional volatility for the next day based on parameters in Table II. Periods end on February 28, 1992, and start in January 1985 (DM), July 1986 (JY), and March 1985 (SF). Regressions use daily observations, and standard errors are corrected for the induced overlap and heteroskedasticity using the Hansen-White (HW) procedure. Asymptotic HW standard errors in parentheses.

Currency	α	Slopes On			R^2
		ISD	MA(20)	GARCH	
DM	0.323*	0.547* [†]			0.1564
	(0.115)	(0.138)			
	0.602*		0.190 [†]		0.0540
	(0.084)		(0.099)		
	0.366			0.478* [†]	0.0499
	(0.191)			(0.227)	
	0.303*	0.669*	-0.099		0.1632
(0.112)	(0.165)	(0.101)			
JY	0.401*	0.622*		-0.173	0.1599
	(0.152)	(0.167)		(0.201)	
	0.327*	0.496* [†]			0.0965
	(0.118)	(0.181)			
	0.563*		0.134 [†]		0.0223
	(0.074)		(0.102)		
	-0.063			1.017*	0.0495
(0.323)			(0.458)		
SF	0.322*	0.578*	-0.073		0.1004
	(0.117)	(0.204)	(0.111)		
	0.042	0.421*		0.474	0.1051
	(0.289)	(0.177)		(0.399)	
	0.392*	0.520* [†]			0.1454
	(0.149)	(0.175)			
	0.658*		0.182 [†]		0.0542
(0.087)		(0.099)			
JY	0.250			0.650*	0.0581
	(0.267)			(0.305)	
	0.370*	0.647*	-0.097		0.1521
	(0.146)	(0.187)	(0.090)		
	0.526*	0.609*		-0.240	0.1490
(0.210)	(0.201)		(0.262)		

* Significantly different from zero at the 5 percent level.

[†] Significantly different from unity at the 5 percent level.

movement in the price S , σ will change by

$$\Delta\sigma = \left(S \frac{\partial\sigma}{\partial c} \right) \frac{\partial c}{\partial S} \left(\frac{\Delta S}{S} \right). \quad (8)$$

For one-month at-the-money options, the delta is about 0.5, and $(\partial c/\partial\sigma)/S$ is about 0.10, using annualized volatility. Now consider that the standard deviation of the S&P 100 Stock Index is about 20 percent, which is 1.3 percent on a daily basis. If some of the 100 stock prices are measured with errors due to stale prices or bid-ask spreads, the error between the measured and true index values could easily reach $\Delta S/S = 0.25$ percent.

Stoll and Whaley (1990), for instance, demonstrate that the S&P 500 Index suffers from serious measurement problems due to infrequent trading and bid/ask price effects. As a result, index prices display very high autocorrelation over 5 to 15 minutes lags, and are led by futures prices. In addition, Harvey and Whaley (1992) report that the stock market and option markets close at 3:00 and 3:15, respectively, and that the index volatility over this 15-minute interval is typically 0.18 percent.

Therefore, using an error of 0.25 percent in the measurement of the underlying spot price appears conservative. Using this number, the error in implied volatility is $(1/0.1) \times 0.5 \times 0.25$ percent = 1.2 percent per annum. This error is very substantial, because it is on the same order of magnitude as the daily variation in actual volatility, and may therefore seriously affect measurements of ISDs. In contrast, these errors are likely to be much smaller for options on currency futures, since there is only one underlying asset, since both the underlying futures and option markets close at the same time, and since both contracts are actively traded. This may explain why implied volatilities are much better predictors of future volatility in the currency futures market than in the S&P 100 Index option market. It would be interesting to check whether ISDs derived from options on S&P 500 futures provide better forecasts than those obtained from the S&P 100 cash options.

V. Simulations

The previous regressions suggest that the ISD is an informative, albeit biased, volatility forecast. Their interpretation, however, is beset by econometric problems. As indicated in Table I, implied and future volatility are highly autocorrelated processes. In such a situation, we could expect to find small-sample biases such as those reported by Dickey and Fuller (1979). When regressing a random variable on its lagged value, they found that the slope coefficient was biased down toward zero in small samples; the bias is generally greatest when the true value of the slope is unity, but also occurs for lower values. The driving factor in their results is the high degree of persistence in the regression variable. Given the persistence of option volatilities, the question is whether this bias also occurs here.

A second problem with the above regressions is that the HW correction is only valid asymptotically. In this case, "large" number of observations is not

measured by the total sample size, which is about 1,800, but rather by the number of truly independent forecasts, which is about 7 years of data multiplied by four contracts per year. Thus we have no measure of how good the asymptotic approximation is.

Finally, and perhaps most importantly, the ISD is still measured with error. Because a closing option quote can represent a bid price, an ask price, or an intermediate price, the observed ISD measures the true ISD with some error. This can be judged from the over the counter market, where foreign currency options are directly quoted in terms of bid/ask volatility, such as "14.2-14.6." Thus a typical spread is 0.4 percent per annum. Conversations with CME traders have confirmed that, for active at-the-money options, spreads are also around 0.3 percent in terms of volatility. Even when averaging over two options, the error drops from 0.4 to 0.28 percent per annum, assuming independent positions within the bid-ask spread. If measurement errors are independent of regression residuals, this errors-in-variables problem will bias the estimated slope coefficient downward. A priori, it is difficult to judge the extent of the bias, since it depends on the relative size of measurement error and of the true time-variation in expected volatility.

To assess whether these problems affect the results, we present simulations that are constructed under the null hypothesis of predictability. We need to design an experiment where the characteristics of volatility closely resemble those of actual data: (i) the future distribution of returns from each day forward must be derived from the present volatility position, (ii) the pseudoimplied volatility must be the average volatility consistent with the underlying volatility process, (iii) each new day resets the whole future pattern of volatility, and the realized volatility is measured from the time-series of ex post returns.

Given that the GARCH(1, 1) model provides an acceptable parsimonious representation of the statistical behavior of volatility, this is the model chosen for the underlying daily volatility. We also keep the actual spacing of contract expiration over the 1985 to 1992 sample period. Since the slope coefficient was the highest for the DM among the currencies examined, the simulations only use DM data.

The pseudo-volatility is measured as follows. Using the actual GARCH parameters in Table II, we generate each day t s conditional volatility over the next day. The variance is also projected for all remaining contract days, using

$$\begin{aligned} h_{t+1} &= \alpha_0 + \alpha_1 r_t^2 + \beta h_t \\ E_t[h_{t+2}] &= \alpha_0 + (\alpha_1 + \beta)h_{t+1}, \\ E_t[h_{t+3}] &= \alpha_0 + (\alpha_1 + \beta)E_t[h_{t+2}], \end{aligned}$$

and so on, which converges to $E_t[h] = \alpha_0 / (1 - \alpha_1 - \beta)$. The average forecast variance on day t is therefore

$$h_{t,T} = (h_{t+1} + E_t[h_{t+2}] + \dots + E_t[h_T]) / (T - t). \quad (9)$$

The following day, we sample from a standard normal distribution¹⁵ the normalized conditional return z_t^* , where the asterisk refers to the fact that this is a pseudorandom variable. The return the following day is generated from $r_t^* = z_t^* \sqrt{h_t^*}$, and so on.

Using the actual contract sequence, a pseudovariance is measured over the remaining life of the contract $\sigma_{t,T}^{*2}$, which is then matched against the average forecast variance. Since the same structure is used to generate the daily innovations and the forecasts, the model is built under the assumption of predictability.

After obtaining as many observations as we had in the original sample, we regress

$$\sigma_{t,T}^* = a + b\sqrt{h_{t,T}^*} + \varepsilon_{t,T}, \quad (10)$$

and record the slope coefficient and associated HW t -statistics. We repeat the experiment 5,000 times, and compare the actual statistics to their empirical distribution, by reporting the "empirical" p -value as the proportion of times the observed statistic was exceeded under the null.

In addition, the simulation experiment can be easily modified to incorporate the effect of errors-in-variables. When regressing $\sigma_{t,T}$ over $\sqrt{h_{t,T}}$, a zero-mean random variable uniformly distributed over the bid-ask spread interval can be added to the regressor.

Table V reports the relevant quantiles, means, and standard deviations of the bootstrapped distributions of regression coefficients and t -statistics under the null hypothesis. With no noise added to the volatility, there is a downward bias in the slope coefficient: its mean and median values are both 0.907, instead of unity. The bias is substantial, since it is on the order of magnitude of one standard error, even with 1,810 observations. The standard deviation of the simulated b s is 0.160, which is only slightly higher than the 0.138 estimate of standard error in Table IV.

The actual slope of 0.547, reported in Table IV, however, falls in the extreme left tail of the empirical distribution, with an empirical p -value of 1.1 percent. Similarly, the HW t -statistic for the hypothesis that $1 - \beta = 0$ is 3.29, which is only exceeded in 1.9 percent of the sample. Alternatively, the empirical distribution of b s shows that the value of 0.547 is $(0.907 - 0.547)/0.160$, or 2.3 standard deviations away from the mean. Although much lower than the t_{HW} of 3.29, this is still significant. Based on this experiment,

¹⁵ Actually, there is evidence that the normalized residuals $r_t/\sqrt{h_t}$ have slightly fatter tails than a normal distribution for daily returns. We also have bootstrapped those from the empirical distribution for the simulations. Because there are more observations in the tails, this lead to distributions of β 's, t 's that have fatter tails than under the normality assumption. Using the bootstrapped sample, we find that the empirical p -value changes from 0.0114 to 0.0384 for the β distribution. Thus rejections still occur when allowing for nonnormal conditional distributions. Note that a misspecification of the distribution function is not a problem here. White (1982) has shown that maximum likelihood estimation using a normal distribution yields consistent estimates of the mean and variance of distributions for which these quantities are finite.

Table V
Simulations of Predictability Regressions Under Null: DM
 Simulated regressions of the future volatility $\sigma_{i,T}^*$ against the volatility forecast $\hat{\sigma}_i^*$:

$$\sigma_{i,T}^* = a + b\hat{\sigma}_i^* + \varepsilon_{i,T}$$

The simulations use 5,000 replications under the null hypothesis of predictability. Future pseudovariances σ^{*2} are computed from daily returns sampled from a distribution following a GARCH(1, 1) process over the life of the contract; average variances are consistent with the underlying GARCH process. Contract expiration cycles are the same as those in the actual sample. The table reports the distribution, mean, and standard error of the b coefficient and of the Hansen-White t -statistics from predictability regressions. The last two columns report the actual statistics as well as the empirical p -value, which is the proportion of times the observed statistic was exceeded under the null. Under Spread = 0.4%, a random error term, distributed uniformly over an interval of 0.4%pa, is added to the regressor; this represents a typical bid-ask spread. Size refers to the number of time-series observations. The full sample corresponds to 1,806 observations, representing data from January 85 to February 92. The shorter sample of 1018 represents 4 years of data.

Statistic	Fractiles of Statistics								Mean	Std. Dev.	Actual Stat.	Empir. p -Value
	0.010	0.050	0.100	0.500	0.900	0.950	0.990	0.990				
Panel A: Spread = 0, Size = 1806												
β	0.541	0.646	0.704	0.907	1.117	1.169	1.281	0.907	0.160	0.547	0.0114	
$t_{HW}(\beta = 0)$	2.888	3.545	3.941	5.481	7.507	8.125	9.653	5.634	1.427	3.966	0.1036	
$t_{HW}(1 - \beta = 0)$	-1.492	-0.916	-0.617	0.575	2.119	2.627	3.707	0.674	1.087	3.290	0.0190	
Panel B: Spread = 0.4%, Size = 1808												
β	0.539	0.640	0.699	0.901	1.108	1.162	1.277	0.901	0.159	0.547	0.0122	
$t_{HW}(\beta = 0)$	2.883	3.546	3.938	5.471	7.485	8.115	9.642	5.625	1.425	3.966	0.1044	
$t_{HW}(1 - \beta = 0)$	-1.441	-0.880	-0.581	0.610	2.159	2.671	3.738	0.712	1.091	3.290	0.0206	
Panel C: Spread = 2.0%, Size = 1806												
β	0.470	0.557	0.611	0.786	0.970	1.017	1.117	0.787	0.140	0.547	0.0420	
$t_{HW}(\beta = 0)$	2.832	3.436	3.812	5.247	7.147	7.777	9.170	5.391	1.348	3.966	0.1316	
$t_{HW}(1 - \beta = 0)$	-0.604	-0.096	0.165	1.410	3.201	3.788	5.106	1.581	1.204	3.290	0.0910	
Panel D: Spread = 0.0%, Size = 1018												
β	0.441	0.562	0.634	0.892	1.153	1.226	1.393	0.894	0.203	0.547	0.0430	
$t_{HW}(\beta = 0)$	2.109	2.591	2.905	4.356	6.541	7.351	9.543	4.601	1.526	3.966	0.3626	
$t_{HW}(1 - \beta = 0)$	-1.528	-0.932	-0.630	0.532	2.295	2.842	4.276	0.703	1.205	3.290	0.0300	

we would conclude that there is some evidence that option volatilities are biased volatility forecasts, assuming that ISDs were measured without error.

The lower panels in the table investigate the impact of measurement errors. When adding noise with mean zero and uniformly distributed on a 0.4 percent per annum range, which is a typical bid-ask spread, the fractiles of the statistics undergo only small downward shifts. The mean of the slope coefficient, for instance, drops to 0.901, and the p -value of the observed statistic increases from 1.1 percent to 1.2 percent. Thus the impact of typical bid-ask spread is small. In order to obtain empirical p -values that approach conventional 5 percent levels, we need to increase the measurement error to 2.0 percent, which is an exceedingly large spread. The third panel in the table shows that the mean then drops to 0.787, with a p -value for the observed β of 4.2 percent, and for the t -statistic of 9.1 percent.

Finally, the lower panel repeats the first experiment with only four years of data, which is similar to the sample size used by Canina and Figlewski (1993). As expected, the distribution of β s is wider than for the whole sample, and empirical p -values now approach the usual 5 percent level. Given the very low slopes observed in their study, however, it is unlikely that the observed statistics would fall well inside the empirical distributions reported here.

Overall, the simulation evidence demonstrates that substantial biases exist in regression tests of the predictive power of ISDs. These biases, however, are unable to explain the observed regression results, unless one accepts rather large measurement errors.

VI. Conclusions

Previous work on the information content of S&P 100 Index options has claimed that ISDs were biased forecasts of future volatility, and were found to be sometimes worse than the simplest statistical models. These results can be given two possible interpretations: either the test procedure is faulty, or option markets are inefficient.

Faulty test procedures can be due to (i) measurement errors, (ii) inappropriate statistical inferences, or (iii) using the wrong pricing model. The option pricing model can be wrong if volatility or interest rates are stochastic, if the value of early exercise is misspecified, or if underlying asset prices do not follow a diffusion process.

We use implied volatilities inverted from the BS model for short-term at-the-money options on the DM, JY, and SF. Previous work has shown that the BS model and stochastic volatility models product ISDs that are quite similar for such options; differences are only on the order of bid-ask spreads.

The present article focuses on measurement errors and statistical inferences as possible explanations for these biases. Because options on currency futures are traded side-by-side with the underlying futures, they are less likely to suffer from the measurement problems affecting cash index options.

We still recognize, however, the effect of bid-ask spreads, and provide simulations to fully account for small-sample biases usually ignored in regression tests of predictive ability. We show that measurement errors can substantially distort inferences.

In contrast with stock index options, we find that statistical time-series models, even when given the advantage of ex post parameter estimates, are outperformed by option-implied forecasts. These results hold for the three currencies under consideration. Even when accounting for possible measurement errors and statistical problems, we still find, however, that ISDs are biased volatility forecasts. The direction of the bias is such that options ISDs appear too variable relative to future volatility.¹⁶ A linear transformation of the ISD provides a superior forecast of exchange rate volatility. These results are consistent with those of Fleming (1993), who analyzes S&P 100 Index options, and also reports that ISDs dominate moving-average volatility, using a more precise estimation technique.

The ultimate test of whether the reported biases are economically, as opposed to statistically, significant is to simulate a dynamic hedging trading rule that attempts to take advantage of mispriced options. The question then is whether such a rule, after transaction costs, generates significant economic profits. This is left for future research.

¹⁶ Overreactions in stock index options have also been reported by Stein (1989).

REFERENCES

- Baillie, R., and T. Bollerslev, 1989, The message in daily exchange rates: A conditional-variance tale, *Journal of Business and Economic Statistics* 7, 297-306.
- Bates, D., and G. Pennacchi, 1991, Estimation of a heteroskedastic state space model of asset prices, Mimeo, University of Pennsylvania.
- Beckers, S., 1981, Standard deviations implied in option prices as predictors of future stock price variability, *Journal of Banking and Finance*, 363-381.
- Black, F., 1976, The pricing of commodity options, *Journal of Financial Economics* 3, 167-179.
- Black, F., and M. Scholes, 1973, The pricing of options and corporate liabilities, *Journal of Political Economy* 81, 637-659.
- Bollerslev, T., 1986, Generalized autoregressive conditional heteroskedasticity, *Journal of Econometrics* 31, 307-327.
- Canina, L., and S. Figlewski, 1993, The informational content of implied volatility, *Review of Financial Studies* 6, 659-681.
- Chesney, M., and L. Scott, 1989, Pricing European currency options: A comparison of the modified Black-Scholes model and a random variance model," *Journal of Financial and Quantitative Analysis* 24, 267-284.
- Chiras, D., and S. Manaster, 1978, The information content of option prices and a test of market efficiency, *Journal of Financial Economics* 10, 29-58.
- Day, T., and C. Lewis, 1992, Stock market volatility and the information content of stock index options, *Journal of Econometrics* 52, 267-287.
- Dickey, D., and W. Fuller, 1979, Distribution of estimators for autoregressive time series with a unit root, *Journal of American Statistical Association* 74, 427-431.
- Duan, J.-C., 1995, The GARCH option pricing model, *Mathematical Finance* 5, 13-32.
- Engle, R., 1982, Autoregressive conditional heteroskedasticity with estimates of the variance of United Kingdom inflation, *Econometrica* 50, 987-1007.

- Fleming, J., 1993, The quality of market volatility forecasts implied by S&P 100 index option prices, Mimeo, Rice University.
- French, D., 1984, The weekend effect and the distribution of stock prices: Implications for option pricing, *Journal of Financial Economics* 13, 547-559.
- Giovannini, A., and P. Jorion, 1989, The time-variation of risk and return in the foreign exchange and stock markets, *Journal of Finance* 44, 307-325.
- Hansen, L., 1982, Large sample properties of generalized method of moments estimators, *Econometrica* 50, 1029-1054.
- , and R. Hodrick, 1980, Forward exchange rates as optimal predictors of future spot rates, *Journal of Political Economy* 88, 829-853.
- Harvey, C., and R. Whaley, 1991, S&P 100 index option volatility, *Journal of Finance* 46, 1551-1561.
- , 1992, Market volatility prediction and the efficiency of the S&P 100 index option market, *Journal of Financial Economics* 31, 43-74.
- Heston, S., 1993, A closed-form solution for options with stochastic volatility with applications to bond and currency options, *The Review of Financial Studies* 6, 327-343.
- Hsieh, D., 1989, Modeling heteroskedasticity in daily foreign exchange rates, *Journal of Business and Economic Statistics* 7, 307-317.
- Hull, J., and A. White, 1987, The pricing of options on assets with stochastic volatility, *Journal of Finance* 42, 281-300.
- Lamoureux, D., and W. Lastrapes, 1993, Forecasting stock-return variance: Toward an understanding of stochastic implied volatilities, *Review of Financial Studies* 6, 293-326.
- Latané, H., and R. Rendleman, 1976, Standard deviation of stock price ratios implied by option premia, *Journal of Finance* 31, 369-382.
- Melino, A., and S. Turnbull, 1990, Pricing foreign currency options with stochastic volatility, *Journal of Econometrics* 45, 239-265.
- Scott, E., and A. Tucker, 1989, Predicting currency return volatility, *Journal of Banking and Finance* 13, 839-851.
- Scott, L., 1988, Random variance option pricing: Empirical tests of the model and delta-sigma hedging, Mimeo, University of Illinois.
- Scott, L., 1987, Option pricing when the variance changes randomly: Theory, estimation and an application, *Journal of Financial and Quantitative Analysis* 22, 419-438.
- Shastri, K., and K. Tandon, 1986, On the use of European models to price American options on foreign currency, *Journal of Futures Market* 6, 93-108.
- Stein, J., 1989, Overreactions in the options market, *Journal of Finance* 44, 1011-1023.
- Stoll, H., and R. Whaley, 1990, The dynamics of stock index and stock index futures returns," *Journal of Financial and Quantitative Analysis* 25, 441-468.
- Whaley, R., 1986, Valuation of American futures options: Theory and empirical tests, *Journal of Finance* 41, 127-150.
- White, H., 1980, A heteroskedastic-consistent covariance matrix estimator and a direct test for heteroskedasticity, *Econometrica* 48, 817-838.
- , 1982, Maximum likelihood estimation of misspecified models, *Econometrica* 48, 817-838.
- Wiggins, J., 1987, Options values under stochastic volatility: Theory and empirical estimates, *Journal of Financial Economics* 19, 351-372.