THE JOURNAL OF FINANCE • VOL. XLIV, NO. 4 • SEPTEMBER 1989

Overreactions in the Options Market

JEREMY STEIN*

ABSTRACT

This paper examines the "term structure" of options' implied volatilities, using data on S&P 100 index options. Because implied volatility is strongly mean reverting, the implied volatility on a longer maturity option should move by less than one percent in response to a one percent move in the implied volatility of a shorter maturity option. Empirically, this elasticity turns out to be larger than suggested by rational expectations theory—long-maturity options tend to "overreact" to changes in the implied volatility of short-maturity options.

ARE INVESTORS "RATIONAL"—DO they behave like Bayesian statisticians—when it comes to incorporating new information into asset prices? Proponents of the efficient markets hypothesis would answer this question in the affirmative. Many others, however, believe that Bayesian rationality is a poor description of investor behavior and that, as a consequence, asset prices tend to be informationally inefficient.

One way in which inefficiencies might arise is through the overreactions of traders to the arrival of new information.[1] This possibility is raised by experimental and survey findings which indicate that people have a systematic tendency to overemphasize recent data at the expense of other information when making projections.[2]

The objective of this paper is to search for evidence of such overreactive behavior in a semingly peculiar place: the options market. Options may appear to be unlikely candidates for overreaction because, unlike primary securities, their prices are closely tied down by arbitrage considerations. As Black and Scholes (1973) demonstrate, if stock price volatility is known, options prices are completely determined—any deviation from the prescribed value implies a riskless profit opportunity.

However, given that volatility is in fact unknown and changing, options do retain something of an independent nature, not wholly redundant with their underlying stocks. Options can be thought of as reflecting a speculative market in volatility—the implied volatility on a given option (obtained by inverting a

* Harvard Business School. I am grateful to Mark Zurack of Goldman Sachs and Co. for providing me with the data on index options and to Bill Fairburn for stellar research assistance. Thanks also to Ken Froot for several helpful discussions and to Fischer Black, René Stulz, and an anonymous referee for their suggestions.

[1] Overreaction to new information is not the *only* way to generate inefficient prices, of course. Purely random noise unrelated to information will cause excess volatility also.

[2] See, for example, Tversky and Kahneman (1974).

Black-Scholes-type formula) should equal the average volatility that is expected to prevail over the life of that option.[3]

This observation suggests that one can conduct "term structure" tests to check whether the temporal structure of options' implied volatilities is consistent with rational expectations. A simple example illustrates the idea behind the tests. Suppose that volatility averages fifteen percent, but that it fluctuates up and down quite rapidly, governed by a strongly mean-reverting process. If a one-month option currently has an implied volatility of twenty-five percent, a two-month option should have an implied volatility that is somewhat lower, with the exact level determined by the coefficient of mean reversion. Conversely, when a one-month option has an implied volatility of five percent, the two-month option should have a higher one.

The evidence presented here contradicts this rational expectations hypothesis for the term structure of implied volatility.[4] Although volatility shocks decay away very quickly, market participants do not take this fully into account when pricing options. Longer term options' implied volatilities move almost in lockstep with those on shorter term options, displaying less of the "smoothing" behavior than is warranted. In this sense, longer term options overreact relative to shorter term ones—they place too much emphasis on innovations in short-term options' implied volatility and too little emphasis on historical data that indicate that these innovations will not persist.

While statistically significant, the mispricings found in the options are small in magnitude compared to those that might conceivably arise with primary securities such as stocks or bonds. Nonetheless, there are a couple of reasons why these mispricings may still be construed as having strong implications. First, option pricing is unique in that it involves only one uncertain variable, namely volatility. In contrast, equity and bond prices contain risk premia. Thus, evidence that appears to suggest irrational excess variability in the stock and bond markets can also be interpreted as rational market responses to time-varying equilibrium rates of return.[5] Since these required rates of return are not observable, it is often hard to come to a definitive conclusion.[6] This ambiguity does not arise in

[3] This is somewhat loose. The statement will be literally true in a world of stochastic volatility only if two conditions are satisfied: 1) there is no risk premium for bearing volatility risk (or, equivalently, volatility stocks are uncorrelated with consumption) and 2) the options pricing formula is linear in volatility, so that there are no Jensen's inequality-type effects. While the latter condition does not hold generally, it does hold almost exactly for at-the-money options, which are the only kind used in the empirical work presented here. These issues are addressed in more detail in Section II.*A*.

[4] More precisely, the evidence rejects the joint hypothesis that the pricing model used is the correct one and that agents form volatility expectations rationally.

[5] One noteworthy example concerns the work of DeBondt and Thaler (1985), who find evidence of a "rebound" effect in the stock market—stocks that had once been extreme losers later tended to perform abnormally well. This can be reconciled with rationality as follows: a decline in stock prices leads to higher leverage, higher CAPM betas, and hence higher expected returns. See DeBondt and Thaler (1987) for a fuller discussion and references.

[6] Time-varying rates of return have also been proposed as an explanation for the mean reversion in stock prices found by Fama and French (1988) and Poterba and Summers (1988).

the context of options, where arbitrage considerations lead to prices that are independent of riskiness.[7]

A second reason why the mispricings found here are noteworthy is that they point to a specific cause of excess fluctuations, that of investor overreaction to new information. This is not the case with much of the "inefficient markets" literature. Furthermore, the inability of traders to behave in a proper Bayesian fashion when pricing longer term options is particularly striking because the problem they face is a relatively simple one: there is one clear source of information, the implied volatility on the short-term option. The only other required inputs are the parameters of the process driving volatility, which are easily objectively quantified in a parsimonious model. Comparing this to the mass of often subjective information that must be sorted through to correctly price stocks, one is tempted to conjecture that simplistic and perhaps overreactive rules of thumb must be of fundamental importance in the stock market.

The remainder of this paper is organized as follows. Section I derives the theoretical relationship between the implied volatilities on options of different maturities that should hold when volatility follows an AR1 process. Section II shows that an AR1 actually does provide a good description of the movements of volatility on S&P 100 options and goes on to test the theoretical relationship of Section I. Section III presents an alternative type of test, one that has the advantage of not requiring any particular specification for the process driving volatility. Section IV offers some conclusions.

I. The Term Structure of Implied Volatility

This section derives the restrictions on relative options prices that are implied by a given stochastic process for stock price volatility. Assume that "instantaneous" volatility σ_t evolves according to the following continuous-time mean-reverting AR1 process:

$$d\sigma_t = -\alpha(\sigma_t - \bar{\sigma})dt + \beta\sigma_t dz. \tag{1}$$

At time t, the expectation of volatility as of time $t + j$ will be given by

$$E_t(\sigma_{t+j}) = \bar{\sigma} + \rho^j(\sigma_t - \bar{\sigma}), \tag{2}$$

where $\rho = e^{-\alpha} < 1$. That is, volatility is expected to decay geometrically back toward its long-run mean level of $\bar{\sigma}$.

Denote by $i_t(T)$ the implied volatility at time t on an option with T remaining until expiration. As noted earlier, this should equal the averaged expected instantaneous volatility over the time span $[t, t + T]$. Using (2), this implies

$$i_t(t) = \frac{1}{T}\int_{j=0}^{T} [\bar{\sigma} + \rho^j(\sigma_t - \bar{\sigma})]\, dj = \bar{\sigma} + \frac{\rho^T - 1}{T \ln \rho}[\sigma_t - \bar{\sigma}]. \tag{3}$$

Equation (3) provides the basis for the empirical tests that follow. It states

[7] Again, this assumes that volatility risk is idiosyncratic. Options prices will always be independent of the risk on the underlying assets.

that, when instantaneous volatility is above its mean level, the implied volatility on an option should be *decreasing* in the time to expiration. Conversely, when instantaneous volatility is below its mean, implied volatility should be increasing in the time to expiration.

Instantaneous volatility σ_t cannot be directly observed. However, this does not interfere with testing of the sort described above. Suppose there are two options: a "nearby," with time to expiration T and implied volatility $i_t^n(T)$, and a "distant," with time to expiration $K > T$ and implied volatility $i_t^d(K)$. Using equation (3), the following restriction can be derived:

$$(i_t^d - \bar{\sigma}) = \frac{T(\rho^K - 1)}{K(\rho^T - 1)}(i_t^n - \bar{\sigma}). \tag{4}$$

Equation (4) is directly testable because it does not involve the instantaneous volatility σ_t. It can be thought of as an elasticity relationship: given a movement in nearby implied volatility i_t^n, there should be a smaller movement in distant implied volatility i_t^d. The exact constant of proportionality depends on the mean reversion parameter ρ, as well as on the times to expiration of the two options.

II. An Empirical Test on S&P 100 Index Options

A. The Data

The data used in this section were generously provided by Mark Zurack of Goldman Sachs and Co. They consist of two daily time series on implied volatilities for S&P 100 index options for the period December 1983 to September 1987, a total of 964 observations for each series. The "nearby" series represents the options with the shortest time to expiration—between zero and one month, depending on the sampling date. The "distant" series represents the options with the next expiration date—between one and two months. Thus, the distant options always have one month longer to go than the nearby options.

For each series, implied volatility is calculated by averaging the implieds on the put and call option closest to being at the money. The options pricing model used to compute implied volatility is a Cox-Ross-Rubinstein (1979) binomial model that explicitly accounts for the dividend yield on the index and for the possibility of early exercise. The inputs are the daily closing prices of the options and the index.

There are several issues worth noting with regard to the quality of the data. First, a study by Evnine and Rudd (1985) found significant evidence of pricing errors in the S&P 100 options market over the period June–August 1984. They attribute some of this to data asynchronies (although they use real-time data) and some to difficulties in arbitraging an entire index. These types of mispricing, which should not lead to a bias in any particular direction, amount to random measurement errors in the implied volatilities being used here. It will be argued below that such measurement errors are not a serious problem for the tests to be performed below. Rather, they have the effect of making the tests *more conservative*—it will be more difficult to reject the null hypothesis of rationality in favor of an overreactions alternative.

A potentially more serious concern would be systematic biases in the implieds due to using an incorrect pricing model. In spite of this paper's premise of stochastic volatility, the pricing model employed in creating the data does not take it into account. Hull and White (1987) and Wiggins (1987) have found that stochastic volatility of the sort assumed here can have a significant impact on some options prices. Fortunately, however, problems are unlikely to arise here because the only data used come from at-the-money options. The prices of such options are almost exactly linear in σ at all maturities. Thus, the implieds derived from inverting a nonstochastic volatility formula should accurately reflect the average volatility that is expected to prevail over the life of these options—there will be no distortions due to using the simpler pricing model.

More fundamentally, it should be emphasized that pricing biases alone would not be enough to cause false acceptance of the overreactions hypothesis in the tests to be performed here. Such false acceptance will occur only if *changes* in the biases are correlated with volatility movements in a very peculiar way (for example, if a rise in volatility systematically increases the extent to which the pricing model overvalues two-month options relative to one-month options). Not only is there no a priori reason for expecting this to be the case but, if biases are small, any daily changes in them are almost certainly of second order.

B. The Persistence of Volatility Stocks

The first step in the empirical analysis is to identify and estimate the stochastic process followed by volatility. Theoretically, we are interested in the serial correlation properties of the instantaneous volatility σ_t, which cannot be observed directly. However, it can be shown that the implied volatility on the nearby option has virtually the same serial correlation characteristics.

Equation (1) specifies the process for volatility as AR1 in levels. In a related study, French, Schwert, and Stambaugh (1987) examine the time-series properties of a measure of volatility constructed directly from stock price data. They find significant skewness in this series and conclude that the stochastic process should be specified in logs. A skewness test was performed here as well, with a different conclusion: there is no significant skewness in implied volatility. This should not be particularly surprising since implied volatilities are not forced to take on huge values after one or two extremely abrupt market moves in the same way as the constructed volatilities of French et al. (1987). In view of the lack of skewness, it seems appropriate to model the volatility process in levels.

The next step is to produce an autocorrelogram and a partial correlogram for i_t^n. Table I presents the results for both the raw i_t^n series and a detrended version. The detrended version is studied because of the presence of a significant ex post trend in volatility, which rises over the sample period. Such a trend can lead to spuriously high autocorrelations when the raw series is examined. In both cases, the autocorrelations have also been converted to an "implied weekly ρ" for ease of comparison. For example, the eight-week autocorrelation is converted to an implied weekly ρ by raising it to the one-eighth power.

The table indicates that an AR1 process provides a good description of the data. A smooth geometric decay in the autocorrelations leads to implied weekly

Table I

Autocorrelograms and Partial Correlograms for i_t^n Series

i_t^n is the implied volatility on the nearby options, and the implied weekly ρ is the autocorrelation raised to the $1/n$ power, where n is the lag length in weeks.

Lag Length	Raw Series			Detrended Series		
(Weeks)	Autocorrelation*	Implied Weekly ρ	Partial Correlation*	Autocorrelation*	Implied Weekly ρ	Partial Correlation*
1	0.906 (0.014)	0.906	0.835 (0.033)	0.874 (0.016)	0.874	0.824 (0.033)
2	0.832 (0.018)	0.912	0.079 (0.043)	0.776 (0.020)	0.881	0.077 (0.043)
3	0.761 (0.028)	0.913	0.054 (0.043)	0.685 (0.024)	0.882	0.051 (0.043)
4	0.685 (0.023)	0.910	−0.089 (0.043)	0.587 (0.026)	0.875	−0.090 (0.043)
5	0.621 (0.025)	0.909	−0.086 (0.043)	0.505 (0.028)	0.872	−0.087 (0.043)
6	0.583 (0.026)	0.914	0.077 (0.043)	0.458 (0.029)	0.878	0.075 (0.043)
7	0.551 (0.026)	0.918	0.014 (0.043)	0.420 (0.030)	0.883	0.013 (0.043)
8	0.530 (0.028)	0.924	0.039 (0.043)	0.394 (0.030)	0.890	0.027 (0.033)

* Standard errors in parentheses.

ρ's that are similar at all lag lengths for both the raw and detrended series (although, as expected, using the detrended series leads to slightly lower estimated autocorrelations). The partial correlations are generally close to zero after the first lag, which is characteristic of an AR1 process. Another argument in favor of the AR1 specification is that adding additional lags does not improve the explanatory power of the regression. For example, with the raw series, the corrected R^2 is 0.82 when only a one-week lag is used, and it is the same when one- *through* eight-week lags are employed.

These conclusions about the appropriateness of an AR1 specification are similar to those reached by Poterba and Summers (1986) in their study of a constructed volatility series. French et al., however, argue that a nonstationary model (specifically, an IMA (1, 3)) fits their data better. In order to check against such a nonstationary alternative, Dickey-Fuller (1981) tests were performed. Using a variety of specifications, the unit root hypothesis is rejected repeatedly, as it was in Poterba and Summers (1986).

Table II recalculates implied weekly ρ's for smaller subsample periods, with each calendar year being treated separately. This is done in order to check that the serial correlation properties of i_t^n are relatively stable. If this were not the case, it would be questionable to use the full sample ρ's as a basis for prediction at any given point in time. The table only displays the annual results for the raw i_t^n series. The detrended series leads to very similar results—the effect of the time trend is less important in shorter sample periods. While the table documents some variation in ρ over time, one conclusion emerges rather strongly: at no point in time is the weekly ρ ever greater than 0.90, and it is often substantially lower, perhaps as low as 0.75.[8]

Before proceeding, two potential problems should be touched on. The first is measurement error. To the extent that the i_t^n's are measured with error, the estimates of ρ given above will be biased downward. However, as will be seen shortly, this bias will be offset in the testing procedure of the next subsection, with the net result being a test that errs on the side of conservatism.

A second possible cause for concern is that, while an AR1 model may be a good simple predictor of future volatility, it need not be the perfect one. Adding perfect further lags or changing the specification could conceivably improve predictive ability somewhat, and it also might be that market participants use information other than current and past volatility when forecasting future levels. In order to ensure that the results of the next subsection are not due to a misspecification of this sort, Section IV develops an alternative type of test, one that does not require any particular specification of how agents forecast future volatility.

C. A Test of Equation (4)

For the data used in this section, equation (4) can be rewritten as

$$(i_t^d - \bar{\sigma}) = \beta(\rho, T)(i_t^n - \bar{\sigma}). \tag{5}$$

[8] Diagnostic tests of the type described for the full sample were also run to see whether the AR1 specification was appropriate for the yearly subsamples. The results were generally similar to those documented in Table I for the full sample.

Table II

Implied Weekly ρ's for Subsample Periods

The implied weekly ρ at a given lag length is the autocorrelation coefficient of nearby implied volatility at that lag, raised to the $1/n$ power, where n is the lag length in weeks.

Lag Length (Weeks)	Full Sample (Raw Series)	Full Sample (Detrended Series)	Annual Periods 1984	1985	(Raw Series) 1986	1987
1	0.906	0.874	0.656	0.874	0.757	0.779
2	0.912	0.881	0.727	0.893	0.737	0.764
3	0.913	0.882	0.748	0.900	0.703	0.731
4	0.910	0.875	0.781	0.898	0.569	0.661

That is, the elasticity β depends on both the decay parameter ρ and the time to expiration T of the nearby option. The time to expiration K of the distant option can be substituted out since for this data set $K = T + 4$ (where time is measured in weeks).[9]

In principle, even for a fixed ρ, equation (5) is nonlinear since the β can be different for different values of T. However, for all practical purposes, this nonlinearity is negligible—over the range of parameters that will be of interest, variations in T have little impact.

Table III illustrates. It gives the values of β for different ρ's and for T's from one to four weeks. As the table shows, when ρ is close to one, β varies very little with changes in T. Indeed, β can be very closely approximated by the following simple function of ρ:

$$\beta \simeq \frac{(1 + \rho^4)}{2}. \tag{6}$$

This function represents a "linear endpoint approximation," with i_t^d given as the average of i_t^n today and the expected value of i_t^n one month hence. The approximation is exact when the nearby option has a maturity T of one month. Otherwise, it is inexact to the extent that volatility does not move linearly over the time period between T and one month.[10] Based on the estimates of ρ in Section II.*B*, the theoretical value of β is somewhere between 0.83 (corresponding to $\rho = 0.90$) and 0.66 (corresponding to $\rho = 0.75$). In other words, if the distant options are priced rationally relative to the nearby options, then, when the nearby implied volatility is one percent above its mean, the distant implied volatility should be *at most* about 0.83 percent above its mean.

One can test to see whether this holds empirically by regressing $(i_t^d - \bar{\sigma})$ against $(i_t^n - \bar{\sigma})$. Before turning to the regression results, the effects of measurement error in the i_t^n's should be discussed. Measurement error has two separate effects:

[9] This is not literally true—in actuality, K will equal either $T + 4$ or $T + 5$, depending on the number of Fridays in the month. However, using $T + 4$ everywhere amounts to a conservative testing procedure—it implies less mean reversion and, hence, will make it more difficult to find overreactions that violate the theoretically correct upper bounds for β.

[10] To the extent that the approximation is incorrect, using it will only make the testing procedure more conservative since the approximation implies mean reversion that is too slow. See footnote 9.

Table III

Values of $\beta(\rho, T) = \dfrac{T(\rho^{T+4} - 1)}{(T + 4)(\rho^{T} - 1)}$

ρ is the autocorrelation coefficient of nearby implied volatility at a one-week lag, and T is the time to expiration of the nearby option. β then represents the theoretically correct elasticity of the distant option's implied volatility with respect to that of the nearby option.

T = No. of Weeks	$\rho = 0.9$	$\rho = 0.8$	$\rho = 0.7$
1	0.82	0.67	0.55
2	0.82	0.68	0.58
3	0.83	0.69	0.60
4	0.83	0.70	0.62
Linear Endpoint Approximation: $\beta = \dfrac{(1 + \rho^4)}{2}$	0.83	0.70	0.62

1) As noted earlier, it leads to a downward bias in ρ^4. Looking at equation (6), this by itself would tend to make us reject too easily the null of rationality in favor of an overreactions alternative—it lowers the "theoretically correct" β and, hence, lowers the empirical β that would have to be found to accept the overreactions hypothesis. 2) On the other hand, measurement error also biases downward the empirical β that will be estimated when i_t^d is regressed against i_t^n. This will work in the opposite direction, making it harder to find overreactions.

Since the coefficient on ρ^4 in equation (6) is ½, the latter effect dominates—the downward bias in our estimated β will be twice the downward bias in the "theoretically correct" β (which is constructed using the biased ρ^4 estimate). Consequently, measurement error will make the test more conservative—it will be more difficult to find evidence of overreactions. This can be seen intuitively by considering the polar case of infinite measurement error. In this case, the estimated value of ρ^4 would be zero, and our "theoretically correct" β would be calculated as ½ from equation (6). However, a regression of i_t^d onto i_t^n would also produce a coefficient of zero, which would look like an *underreaction* relative to the "theoretically correct" regression coefficient of ½.

Table IV presents the results of the regression tests, both for the full sample and for each year run separately. All the regressions are OLS, and the standard errors have been corrected using the methodology of Hansen (1982) to deal with serial correlation in the residuals.[11] The coefficients for the full sample and for the first two years are significantly higher than the highest plausible theoretical β of 0.82. The coefficients for the last two years are somewhat lower. However, if one refers back to Table II and notes that there is no evidence of a ρ above 0.8 (which implies a theoretical β of 0.70) during these two years, they appear to be characterized by overreactions as well; while the elasticity of i_t^d with respect to

[11] The regressions use "raw" i_t^d and i_t^n data. However, using detrended data leads to virtually identical results. For example, the full sample coefficient using detrended data is 0.922, as opposed to the value of 0.929 reported here for the raw data.

Table IV

Regressions of i_t^d onto i_t^n

i_t^d is the implied volatility on the distant option series, and i_t^n is the implied volatility on the nearby option series.

Sample Period	Coefficient	Standard Error	R^2
Full Sample	0.929	(0.012)	0.96
1984	0.904	(0.042)	0.81
1985	0.940	(0.042)	0.92
1986	0.818	(0.028)	0.92
1987	0.819	(0.023)	0.94

i_t^n is lower, there is also a concurrent drop from the previous year in the persistence of volatility shocks.

It bears repeating that measurement error in the i_t^n's will tend to bias downward the regression coefficients and, on net, make the tests more conservative. Further downward bias might arise if the daily i_t^d and i_t^n data are not perfectly synchronized. To investigate this possibility, the regressions were rerun using weekly averages of the daily data. As it turned out, this had little effect on the results, suggesting that asynchronies of this type are not a problem.

III. An Alternative Test for Overreactions

The tests of the previous section relied on specifying a particular stochastic process for volatility. This section considers a more general test procedure which makes no such requirement. It is exactly analogous to the more recent tests of rational expectations in the term structure of interest rates, such as those performed by Campbell and Shiller (1984), among others. (See Froot (1989) for a survey of much of this literature.)

The linear endpoint approximation of equations (5) and (6) can be rewritten in more general form as

$$(i_t^d - \bar{\sigma}) = ½[(i_t^n - \bar{\sigma}) + E_t(i_{t+4}^n - \bar{\sigma})]. \tag{7}$$

This in turn can be rearranged to yield

$$E_t[(i_{t+4}^n - i_t^n) - 2(i_t^d - i_t^n)] = 0. \tag{8}$$

The prediction error given within the expectations operator on the left-hand side of (8) should, according to rational expectations, be white noise. If, on the other hand, there are overreactions, this prediction error will be negatively correlated with i_t^n—when i_t^n is high, i_t^d is *too high,* leading to a negative prediction error.

Thus, one can test for overreactions simply by regressing the prediction error on i_t^n. Table V describes the results. Again, the regressions are OLS and the standard errors have been corrected to account for serial correlation of the residuals induced by the overlapping observations. The results generally bear out the overreactions hypothesis. The coefficients for the full sample and for each individual year are all negative. In all but one instance, the coefficients are either

Table V

Regressions of $[(i^n_{t+4} - i^n_t) - 2(i^d_t - i^n_t)]$ onto i^n_t

i^d_d is the implied volatility on the distant option series; i^n_t is the implied volatility on the nearby option series; and i^n_{t+4} is i^n_t led by four weeks.

Sample Period	Coefficient	Standard Error	t-Statistic
Full Sample	−0.166	(0.093)	−1.78
1984	−0.602	(0.150)	−4.01
1985	−0.346	(0.248)	−1.40
1986	−0.611	(0.206)	−2.97
1987	−0.461	(0.259)	−1.78

significantly different from zero at the five percent confidence level or very close to it.[12]

The point estimates in Table V are roughly consistent with the magnitudes seen in the previous section. Suppose that the theoretically correct β is 0.70 (corresponding to a weekly ρ of 0.80) but that the empirically measured elasticity is 0.90. Substitution into equation (8) yields a slope of −0.40 for the residual with respect to i^n_t, which is in the range seen in the annual regressions of Table V.

IV. Conclusions

Two types of tests have been presented, both of which reject the joint hypothesis that the pricing model used to recover implied volatilities is correct and that volatility expectations are formed rationally. If one maintains the assumption that the pricing model is correct, the tests are evidence of consistent overreactive behavior in the term structure of options' implied volatilities. The first test used an AR1 model for volatility to derive theoretical upper bounds for the elasticity of distant implied volatility with respect to nearby implied volatility. It was then found that the empirical values of this elasticity exceeded the theoretical upper bounds.

The second test dispensed with the need for a specific stochastic model of volatility. Instead, it employed a more conventional rational expectations procedure, checking to see whether white noise residuals resulted when the $(i^d_t - i^n_t)$ spread was used to forecast future volatility changes. The negative coefficients obtained when the residuals were regressed on i^n_t reinforced the conclusions drawn from the first test.

In spite of their statistical significance, the overreactions found here are fairly

[12] A couple of qualifying points should be made with respect to the results of Table V. First, not all months contain four weeks—some contain five. However, the same type of argument as in footnote 9 implies that using i^n_{t+4} instead of i^n_{t+5} only makes it more difficult to find evidence of overreactions. (This is supported by rerunning the regressions with i^n_{t+5}, which leads to stronger results.) Second, as discussed in the last section, the type of linear endpoint approximation used in equation (7) is inexact when T is less than four weeks (although for reasonable parameter values, it is very accurate). Again, however, an argument similar to that in footnote 10 establishes that using the approximation simply amounts to a conservative testing procedure.

small in economic magnitude. For example, if the theoretically correct elasticity is 0.85 but the actual is 0.95, then the distant options have an implied volatility that is typically one percent off the mark for every ten percent that volatility is away from its mean. Suppose that volatility averages fifteen percent, but that the implied on the nearby is currently at twenty-five percent. Then the implied on the distant "should" be at 23.5 percent (23.5 = 15 + 0.85 (25 − 15)) but will tend to be too high at around 24.5 percent. For a two-month European call option at the money, with an interest rate of ten percent and an initial stock price of 100, the "correct" volatility of 23.5 percent gives a price of 4.66. The "incorrect" volatility of 24.5 percent gives a price of 4.82—an overpricing of 3.4 percent. Even a very "aggressive" example leads to modest mispricings. If the theoretically correct elasticity is only 0.65, the correct volatility of 21.5 percent yields a price of 4.35, so that the incorrect price of 4.82 represents a 10.8 percent overpricing.

Of course, even the most naive rule of thumb cannot have all that large an impact if we are only comparing one- and two-month options. Suppose that investors ignored mean reversion in volatility altogether, so that one- and two-month options always had the same implied volatilities. The same reasoning as above shows that, once again, the deviations from theoretically correct prices would be small.

In order for such ignorance of mean reversion in volatility to translate into large pricing errors, the options involved would have to be quite long-lived. Thus, it would be very interesting to know whether investors in long-dated options also tend to overreact to changes in short-term implieds. If they do, the mispricings involved are likely to be much more economically significant.

Because of this potential interest, an attempt was made to construct tests analogous to those presented above, using data on warrants and short-term options for several individual companies that had both available. For example, American General has outstanding warrants that expire in 1989, as well as short-term options. The prices of these securities were used to create the data series i_t^d (the implied on the warrant) and i_t^n (the implied on the option) for American General. Regressions were then run to estimate the elasticity of the former with respect to the latter. Unfortunately, the data series were simply not of sufficient quality to generate any meaningful inferences. Unlike with the S&P 100 data, the implieds contain very large measurement errors due to infrequent trading and other problems. This was the case with all of the six companies studied.[13]

In spite of the lack of success in this endeavor, the methodology may yet prove useful in studying longer dated options. If, as is likely, markets for long-dated index options develop over the next few years, they may provide a high-quality source of data more suitable for the type of tests developed in this paper.

[13] In addition to American General, these companies were Asarco, Apache, FNMA, Lilly, and Southwest Air.

REFERENCES

Black, Fischer and Myron Scholes, 1973, The pricing of options and corporate liabilities, *Journal of Political Economy* 81, 637–654.

Campbell, John and Robert Shiller, 1984, A simple account of the behavior of long-term interest rates, *American Economic Review Papers and Proceedings* 74, 44–48.

Cox, John, Stephen Ross, and Mark Rubinstein, 1979, Option pricing: A simplified approach, *Journal of Financial Economics* 7, 229–263.

DeBondt, Werner and Richard Thaler, 1985, Does the stock market overreact?, *Journal of Finance* 40, 793–805.

——— and Richard Thaler, 1987, Further evidence on investor overreaction and stock market seasonality, *Journal of Finance* 42, 557–581.

Dickey, David and Wayne Fuller, 1981, Likelihood ratio statistics for autoregressive time series with a unit root, *Econometrica* 49, 1057–1072.

Evnine, Jeremy and Andrew Rudd, 1985, Index options: The early evidence, *Journal of Finance* 40, 743–755.

Fama, Eugene and Kenneth French, 1988, Permanent and temporary components of stock prices, *Journal of Political Economy* 96, 246–273.

French, Kenneth, William Schwert, and Robert Stambaugh, 1987, Expected stock returns and volatility, *Journal of Financial Economics* 19, 3–29.

Froot, Kenneth, 1989, New hope for the expectations hypothesis of the term structure of interest rates, *Journal of Finance* 44, 283–305.

Hansen, Lars Peter, 1982, Large sample properties of generalized method of moments estimators, *Econometrica* 50, 1029–1054.

Hull, John and Alan White, 1987, The pricing of options on assets with stochastic volatilities, *Journal of Finance* 42, 281–300.

Poterba, James and Lawrence Summers, 1986, The persistence of volatility and stock market fluctuations, *American Economic Review* 76, 1142–1151.

——— and Lawrence Summers, 1988, Mean reversion in stock prices: Evidence and implications, *Journal of Financial Economics* 22, 27–59.

Tversky, Amos and Daniel Kahneman, 1974, Judgement under uncertainty: Heuristics and biases, *Science* 185, 1124–1131.

Wiggins, James, 1987, Options values under stochastic volatility, *Journal of Financial Economics* 19, 351–372.