

Implied volatility functions: a reprise

July 1999

Joshua V. Rosenberg
Department of Finance
NYU - Stern School of Business
44 West 4th Street, Suite 9-190
New York, New York 10012-1126
(212) 998-0311
jrosenb0@stern.nyu.edu

Abstract

Dumas, Fleming, Whaley (DFW, 1998) find that option models based on deterministic volatility functions (DVF) perform poorly because the estimated volatility function is unstable over time. DFW provide evidence that the DVF changes significantly on a weekly basis.

This paper proposes a new class of dynamic implied volatility function models (DIVF). This class of models separates a time-invariant implied volatility function from the stochastic state variables that drive changes in the individual implied volatilities. The dynamics of the state variables are modeled explicitly. This framework facilitates consistent pricing and hedging with time-variation in the implied volatility function (IVF).

In tests conducted using the full history of S&P500 futures option prices, the DIVF model is found to substantially improve pricing performance compared to static implied volatility function models and benchmark pricing models such as Black and Scholes (1973).

This paper has benefited from the suggestions of David Bates and David Backus. The most recent version of this paper is available at http://www.stern.nyu.edu/~jrosenb0/J_wpaper.htm. ©1999.

I. Introduction

Dumas, Fleming, Whaley (DFW, 1998) find that the class of deterministic volatility models (DVF), which is based on approximating a diffusion in which the (local) underlying asset volatility is a deterministic function of the underlying price, perform poorly because this function changes significantly on a weekly basis. DFW also test an implied volatility function (IVF) model which is found to outperform the DVF models, but still exhibits instability. Similar results are found in Jackwerth and Rubinstein (1998).

DVF models were hoped to offer an improvement over previous techniques based on an implied volatility function (IVF) which relates option specific implied volatilities to option specific characteristics. While IVF models are adequate for interpolative pricing, these models have been considered to be limited by the undefined relationship between the risk-neutral process and the IVF. Recently, Reiner (1999) has solved this problem for implied volatility functions defined in terms of the option exercise price and proportional moneyness.

Some of the criticism of the ad-hoc nature of IVF models is misguided, since these models do not assume that the Black-Scholes model is correct. Instead, they take the Black-Scholes pricing function as a useful tool to normalize option prices along a common scale of implied volatilities. The market pricing function is then defined in terms of the implied volatilities. As long as the estimated implied volatility function is time-invariant, these models are potentially attractive for option pricing and hedging.

However, a second limitation of IVF models, which was hoped to be overcome using DVF models, is the time-instability of the estimated implied volatility function. In common use, implementation of IVF models requires re-estimation the IVF on a daily or weekly basis to remedy the observed functional instability. It is this practice which is ad-hoc and inconsistent, since the IVF is allowed to exhibit dynamics which are outside of the static specification. The hedge ratios which are consistent with the ad-hoc implementation of IVF models are undetermined, since the relationship between changes in the implied volatility function and other state variables is unspecified.

Since the DVF approach has not been found to be successful, it may be useful to re-examine the IVF approach and consider generalizations which might remedy the problem of time-instability. For example, one limitation of the DVF technique — not necessarily shared by the IVF technique — is the restrictive formulation of the volatility function, which only incorporates the current underlying price (S_t) and time (t) as state variables.

(1) DVF: $\sigma_t = \sigma(S_t, t)$

In contrast, IVF models allow implied volatilities to depend on a wide variety of possible state variables. In the existing literature, IVF models utilize contract specific characteristics as state variables. For example, the implied volatility function is defined over option exercise prices (K) in Shimko (1993) and Bates (1995). In another common formulation, the implied volatility function is defined over proportional moneyness ($K / S_t - 1$). For example, this is used in Skiadopoulos, Hodges, and Clewlow (1999) in a principal components analysis of properties of the implied volatility surface. In this paper, option contracts are aggregated over maturity ranges to adjust for maturity effects.

Natenberg (1994) suggests that implied volatility function stability is improved by adjusting proportional moneyness for maturity. Tompkins (1995) includes a further refinement which defines standardized moneyness as the ratio of maturity-adjusted-proportional-moneyness and at-the-money implied volatility ($\sigma_{ATM,t}$). In this case, moneyness is measured in standard deviation units.

In order to be internally consistent, implied volatility function models require that the volatility function is time-invariant. A general formulation of the IVF model, with $\sigma_{i,t}$ as the date t implied volatility of an option with exercise price K_i , is:

(2) IVF: $\sigma_{i,t} = f(S_t, K_i, T-t, \sigma_{ATM,t})$

This paper proposes and investigates a class of dynamic implied volatility function models (DIVF). This class of models separates the time-invariant implied volatility function from the stochastic state variables which drive changes in the individual implied volatilities. The dynamics of the state variables are modeled explicitly. This framework facilitates consistent pricing and hedging with time-variation in the IVF.

This is a potentially promising approach, because, as noted in DFW, a central reason for the failure of the DVF models is that they are unable to replicate the stylized fact that “New market information induces a shift in the level of overall market volatility from week to week (DFW, p. 2081).” An IVF model which incorporates this feature may be able to extract the useful information embedded in the cross-section of option prices, while preserving the key stochastics that drive changes in individual option implied volatilities.

To this end, the DIVF is defined so that each option's implied volatility depends on the level of the at-the-money implied volatility. This is accomplished by identifying a time-invariant function of the relative implied volatility of each option (with exercise price K_i) expressed as $\sigma_{rel,i,t} = \sigma_{i,t} / \sigma_{ATM,t} - 1$ in terms of the contract specific characteristics (Equation 3). The at-the-money implied volatility is treated as the stochastic driving variable for IVF dynamics. To incorporate empirical evidence of volatility mean-reversion and asymmetric effects of returns on volatility, the ATM implied volatility is represented as function of lagged ATM implied volatility ($\sigma_{ATM,t-1}$) and the lagged asset return (r_{t-1}) (Equation 4). Equation (5) defines the contract specific implied volatility in terms of the relative implied volatility and the at-the-money implied volatility.

$$(3) \quad \text{DIVF: } \sigma_{rel,i,t} = f(S_t, K_i, T-t, \sigma_{ATM,t})$$

$$(4) \quad \text{DIVF: } \sigma_{ATM,t} = g(\sigma_{ATM,t-1}, r_t)$$

$$(5) \quad \text{DIVF: } \sigma_{i,t} = \sigma_{ATM,t}(\sigma_{rel,i,t} + 1)$$

Empirical results using the full history of S&P500 futures options indicate that a DIVF model can generate substantially improved pricing performance over static implied volatility function models and benchmark pricing models such as Black and Scholes (1973). This is strong evidence indicating that the pricing formula should include a stochastic state variable which scales the level of the implied volatility function.

This paper is related to Britten-Jones and Neuberger (1999) which describes a theoretical framework for estimation of risk-neutral processes consistent with stochastic volatility. In contrast to their paper which suggests use of an implied generalized trinomial lattice as an approximating process, this paper selects a particular discrete-time model which incorporates key features of the dynamics of implied volatility functions and tests the estimated model using data from a traded options market.

The remainder of this paper is structured as follows. Section II describes the specification of the DIVF model, and estimates the DIVF model using S&P500 futures option data. Section III presents pricing and hedging tests. Section IV concludes the paper.

II. The dynamic implied volatility function model

The DIVF used for estimation in this paper utilizes the Tompkins (1995) definition of moneyness in standard deviations, which will be referred to as standardized moneyness. This definition has the advantage of normalizing moneyness for time to expiration and volatility simultaneously. However, the adequacy of this specification in terms of its time-invariance is ultimately an empirical question.

$$(6) \quad \text{DIVF: } \sigma_{i,t}(S_t, K_i, T-t, \sigma_{\text{ATM},t}) = f[(S_t / K_i - 1) / (\sigma_{\text{ATM},t} \sqrt{T-t})]$$

Using an aggregated time-series and cross-section of option implied volatilities, the function $f[\bullet]$ may be estimated non-parametrically using a kernel (e.g. Ait-Sahalia and Lo (1998)), semi-parametrically using splines (e.g. Rosenberg and Engle (1999)), or parametrically using an approximating polynomial function (e.g. Shimko (1993)).

The dynamics of the implied volatility function are incorporated through an explicit specification of the at-the-money implied volatility process. The current at-the-money implied volatility is modeled as a function of its first lag (to incorporate volatility mean-reversion) and the most recent asset return (to incorporate asymmetric return effects). The formulation of the asymmetry term is related to that of Glosten-Runkle-Jagannathan (1993); this specification was found to exhibit superior performance in objective volatility modeling by Engle and Ng (1993).

$$(7) \quad \text{DIVF: } \sigma_{\text{ATM},t} = \omega + \alpha \sigma_{\text{ATM},t-1} + \gamma |\max(0, -r_t)|$$

There are two differences in this model from the GJR formulation. First, the model is written in standard deviation rather than variance terms. Second, the timing appears different in this formulation than in GJR, although it is analogous to GJR. In this model, the timing of the measurements is at the end of each discrete period. At date t , the end-of-day implied volatility is represented as a function of the observed previous end-of-day implied volatility and the return over the day ($r_t = S_t/S_{t-1} - 1$). If the at-the-money implied volatility is viewed as the conditional volatility over the subsequent period (i.e. $\sigma_{\text{ATM},t-1} = \sigma_{t|t-1}$ and $\sigma_{\text{ATM},t} = \sigma_{t+1|t}$, see, e.g.) then this model has the same timing as the GJR model.

The parameters of equation (7) are most easily estimated by appending a white noise error term to the equation and using linear regression. Note that the returns process does not need to be specified for continuous hedging; the hedge ratios will depend on the ω , α , and γ parameters which may be estimated using historical data. To construct a minimum variance hedge, an additional equation describing the

underlying asset return process would be required. In this case, estimation would be accomplished using data of the same frequency as the rebalancing interval.

Equations (5) - (7) define an implied volatility function which is stochastic, with dynamics driven by underlying asset returns. In contrast, DVF models require that the local volatility is a function of the asset price level. Previous (static) IVF models require that implied volatilities are a constant function of exercise price or moneyness.

Option hedge parameters for IVF and DIVF models are obtained by differentiating the Black-Scholes pricing formula, with the volatility parameter replaced by the implied volatility function, once and twice with respect to the current underlying price (S_t). Let $\Delta_{BS}(\bullet)$, $\Gamma_{BS}(\bullet)$, and $\Lambda_{BS}(\bullet)$ represent the Black-Scholes delta, gamma, and vega hedge parameters. Then:

$$(8) \quad \Delta_{(D)IVF}(K_i) = \frac{\mathbb{1}C_t(S_t, \mathbf{s}_{i,t}(S_t))}{\mathbb{1}S_t} = \Delta_{BS}(\bullet) + \Lambda_{BS}(\bullet) \frac{\mathbb{1}\mathbf{s}_{i,t}(S_t)}{\mathbb{1}S_t}$$

$$(9) \quad \Gamma_{(D)IVF}(K_i) = \frac{\mathbb{1}^2 C_t(S_t, \mathbf{s}_{i,t}(S_t))}{\mathbb{1}S_t^2} = \Gamma_{BS}(\bullet) + \Lambda_{BS}(\bullet) \frac{\mathbb{1}^2 \mathbf{s}_{i,t}(S_t)}{\mathbb{1}S_t^2} + \dots$$

Hedge parameters for IVF and DIVF models differ in the specification of the volatility function and the resulting derivatives. In static IVF models, an underlying price change effects the implied volatility through a change in the option moneyness (equation 3). In the DIVF model, an underlying price change the implied volatility through the at-the-money implied volatility as well as the option's relative moneyness (equations 5-6).

The cumulative risk-neutral density function for IVF and DVF models is derived by differentiating the pricing formula once with respect to the exercise price, as noted by Breeden and Litzenberger (1978). The risk-neutral density function is obtained by differentiating twice. For the DVF model, this formula provides the risk-neutral density for all maturities, even when estimation is conducted using options with the same maturity. Let $N(\bullet)$ be the standard cumulative normal density function and $d_2 = [\ln(S_t/K_i) + 1/2(r - \sigma_{i,t}^2)(T-t)] / \sigma_{i,t}\sqrt{T-t}$. Then:

$$(10) \quad F_t^*(S_T | S_T \leq K_i) = e^{r(T-t)} \frac{\mathbb{1}C_t(K_i, \mathbf{s}_{i,t}(K_i))}{\mathbb{1}K_i} = (1 - N(d_2)) + e^{r(T-t)} \Lambda_{BS}(\bullet) \frac{\mathbb{1}\mathbf{s}_{i,t}(K_i)}{\mathbb{1}K_i}$$

$$(11) \quad \frac{\mathbb{1} \mathbf{s}_{i,t}(K_i)}{\mathbb{1} K_i} = \frac{\mathbb{1} \mathbf{s}_{i,t}(K_i)}{\mathbb{1} M_i} \frac{\mathbb{1} M_i}{\mathbb{1} K_i} = \frac{\mathbb{1} \mathbf{s}_{i,t}(K_i)}{\mathbb{1} M_i} \frac{1}{S_t \mathbf{s}_{ATM,t} \sqrt{T-t}}$$

$$(12) \quad M_i = \frac{K_i / S_i - 1}{\mathbf{s}_{ATM,t} \sqrt{T-t}}$$

III. Testing the DIVF model

III.i. IVF, DIVF, and ad-hoc models

One of the key findings of DFW (1998) is that deterministic volatility functions are not stable over time. This results in poor pricing and hedging performance. In this section, the stability and pricing performance of static implied volatility functions and dynamic implied volatility functions are compared.

One stability test used in DFW “evaluates how well each week’s estimated volatility function values the same options one week later (DFW, p. 2081).” In fact, this test is biased favorably towards DVF models, since it allows weekly time-variation in the DVF. A test which is consistent with the DVF specification would require the DVF to be constant over the entire sample period. This paper considers both types of tests.

Three IVF models are estimated maintaining the assumption that the IVF is time-invariant. The first model has an implied volatility function equal to a scalar which is the average at-the-money implied volatility over the sample period. This is consistent with the Black-Scholes (1973) model. The second model has an implied volatility function which depends on the contract specific exercise price. The third model has an implied volatility function which depends on the contract specific moneyness. The fourth estimated model is the DIVF model described in the previous section.

Model 1: IVF (Black-Scholes)

$$\mathbf{s}_{i,t} = \bar{\mathbf{s}}_{ATM}$$

Model 2: IVF (Exercise price)

$$\mathbf{s}_{i,t} = f(K_i)$$

Model 3: IVF (Proportional moneyness)

$$\mathbf{s}_{i,t} = f(K_i / S_t - 1)$$

Model 4: DIVF (Rel. imp. vol., std’ized m’ness)

$$\mathbf{s}_{i,t} = \mathbf{s}_{ATM,t} [f((K_i / S_t - 1) / \mathbf{s}_{ATM,t} \sqrt{T-t}) + 1]$$

Three ad-hoc (AH) alternative IVF models are included for comparison. These models allow weekly time-variation in the IVF (denoted by $f_{t-1}(\bullet)$), and are comparable to those tested by DFW. However, these models are not internally consistent, since IVF models utilize a static specification. The models are suited to out-of-sample testing, because they model current implied volatilities using a function estimated using the prior week's data.

Model 5: Ad-hoc IVF (Black-Scholes)

$$\mathbf{s}_{i,t} = \mathbf{s}_{ATM,t-1}$$

Model 6: Ad-hoc IVF (Exercise price)

$$\mathbf{s}_{i,t} = f_{t-1}(K_i)$$

Model 7: Ad-hoc IVF (Proportional moneyness)

$$\mathbf{s}_{i,t} = f_{t-1}(K_i / S_t - 1)$$

III.ii. Model estimation

This paper utilizes the full history of S&P500 futures option prices traded on the Chicago Mercantile Exchange for estimation of the IVF and DIVF models. A daily record of all contracts traded along with settlement price, volume, and other statistics was obtained from the Futures Industry Institute covering the period from 1983 through 1998. In addition, a similar database with daily settlement prices for the corresponding S&P500 futures contracts was obtained. Other data used includes S&P500 index levels and returns obtained from the CRSP database and BBA three-month LIBOR rates from the Datastream database, which are used to proxy for the discount rate.

For this study, S&P500 futures options with the March expiration cycle are utilized, since these options have the same expiration as the corresponding futures contract. An advantage of the use of futures options versus index options is that the daily futures options database contains settlement prices which reflect market conditions at the close of trading for each contract. Under most circumstances, the settlement price is the average of the highest and lowest transaction prices in the last 30 seconds of trading. Additional rules are used for options that do not trade at the end of the day; see Rule 813 of the Chicago Mercantile Exchange Rulebook. This is in contrast to the daily CBOE S&P500 index options database which is potentially subject to problems with price synchronicity. The reported S&P500 closing option prices from the CBOE are the last trade of the day, which might occur in the morning.

Implied volatilities for each option are calculated by numerically inverting the Black (1976) formula. While S&P500 futures option contracts have American exercise-style, a comparison of implied volatilities estimated using the Barone-Adesi Whaley (1987) model — which incorporates the value of early exercise — and the Black model shows that differences are negligible. The call and put implied

variances for contracts with identical exercise price and maturity are averaged to obtain a single implied volatility estimate for each exercise price.

The following data screening conditions are used to minimize the influence of outlying and erroneous observations. First, futures option contracts with little liquidity, reflected by fewer than 5 contracts traded over the day, are excluded from the analysis. Second, contracts for which there are apparent data errors, reflected by annualized implied volatilities less than 1% or greater than 90% are excluded from the analysis. Third, short and medium term expirations are included in the analysis using the criterion that the number of calendar days until expiration is between 10 and 100. Fourth, exceptionally deep in-the-money and out-of-the-money options (i.e. absolute standardized moneyness greater than five standard deviations or absolute proportional moneyness greater than 25%) are excluded. Finally, days for which there are fewer than four options meeting the above criteria are excluded.

Table 1 summarizes the data showing that over the sample period, there are 3,140 trading days with 53,714 contracts meeting the screening criteria. Figure 1 plots the at-the-money implied volatility on each Wednesday of the sample, showing that it varies considerably over the period. Figure 1 also plots the fitted at-the-money implied volatility using the implied volatility functions estimated for models 2 and 3 with the option exercise price set equal to the contemporaneous futures price. The fitted at-the-money implied volatility for model 4, using equation (7) is included as well.

The behavior of fitted implied volatilities for at-the-money contracts varies widely across the different models. Following the behavior of the at-the-money contract is equivalent to following the behavior of a contract with exercise price equal to the contemporaneous futures price. In the proportional moneyness model, this results in a time-invariant at-the-money implied volatility equal to the model intercept, since the proportional moneyness is equal to zero.

Figure 1 reveals that models 1 through 3 fail to characterize many features of at-the-money volatility; a similar result is noted by DFW (p. 2079). Hence, an improved model should incorporate a separate stochastic factor for the volatility level. Model 4 incorporates the time-series properties of at-the-money implied volatility and provides a good fit to the actual at-the-money implied volatility.

Figures 2 through 4 plot the implied volatility function relationships for models 2 through 4. Also plotted is approximating cubic polynomial function which summarizes this relationship. Evidence for time instability of the implied volatility function defined over exercise prices is clear from Figure 2. As the index level rises, the implied volatility curve shifts to the right, resulting in a sequence of negatively sloped curves. Similarly, evidence for time instability of the implied volatility function defined over proportional moneyness is clear from Figure 3. When at-the-money volatility increases, the volatility

curve shifts upwards, resulting in a stack of similarly shaped curves. The strong relationship between relative implied volatilities and standardized moneyness is shown in Figure 4.

The first panel of Table 2 reports measures of fit for the cubic approximation (estimated by polynomial regression) for each model and properties of the fitted implied volatilities. The adjusted- R^2 from the regression measures the explanatory power of the implied volatility function. Model 2 has an adjusted- R^2 of 11% versus 40% for model 3 and 91% for model 4.

These results may also be interpreted as a test of the stability of the implied volatility function. Low adjusted- R^2 values suggest a great deal of implied volatility variation occurs which is not explained by a time-invariant function of exercise price or proportional moneyness. The dynamic specification, which incorporates time-variation in the at-the-money implied volatility, is clearly superior although there is some variability left unexplained.

The second panel of Table 2 reports estimates for models 5 - 7. These models are estimated each Wednesday of the sample period (647 times) and average parameter values are reported. The difference in fit for the weekly re-estimation is dramatic with models 6 and 7 (ad-hoc versions of models 2 and 3) having average adjusted r-squared values of about 98%. The average parameter values are also quite different than for the full sample estimation. This is also evidence of functional time-instability.

The third panel of Table 2 reports estimates for the at-the-money implied volatility process used in the DIVF model (equation 7). Estimation is performed using data from each Wednesday of the sample period. The Wednesday at-the-money implied volatility is regressed on its one week lag and the GJR function of the one-week S&P500 index return. This model effectively characterizes S&P500 at-the-money volatility dynamics with an adjusted r-squared of 90%. As expected, lagged implied volatilities are helpful in predicting future implied volatilities; the α parameter estimate is .89. In addition, the leverage effect is apparent. A negative return results in an increase in at-the-money implied volatility; the β parameter estimate is .06.

The fourth panel of Table 4 reports characteristics of the fitted implied volatilities obtained using the seven estimated models. A lower bound of 1% annualized volatility is used for models 6 and 7. Some of the models exhibit bias, as evidenced by the difference between average fitted and actual implied volatility. IVF models 1 - 3 are unable to match the observed variability of implied volatilities across time and moneyness.

III.iii. Model testing

DFW propose and implement a test to measure stability and pricing performance of DVF models based on the ability of these models to fit observed option prices. This paper utilizes the price difference criterion proposed in DFW, but also considers an implied volatility difference criterion, which normalizes results across option moneyness and maturity. This paper also utilizes a significantly longer testing period (15½ years) than DFW (5½ years).

Following DFW, the cross-section of actual option prices is extracted each Wednesday (10,855 observations), and contract specific implied volatilities are calculated using the Black (1976) model. Models 1 through 7, estimated in Section III.ii., are then used to estimate fitted implied volatilities. Fitted put option prices are obtained using the Black (1976) model, the fitted implied volatilities, the contemporaneous futures prices and riskless rates, and the contract specific maturities and exercise prices.

DIVF implied volatilities are obtained using the estimated relative implied volatility function (equation 6). The implied volatilities are calculated using equation (5) and the fitted at-the-money implied volatilities using equation (7).

Table 3 reports measures of pricing accuracy for each model. The first panel reports implied volatility errors, defined as the observed implied volatility minus the fitted implied volatility. The second panel reports pricing errors, defined as the observed price minus the fitted price.

The test results are consistent across comparisons of pricing and implied volatility errors. The DIVF model provides the lowest pricing error standard deviation of any of the models, including the three ad-hoc models. The difference is most dramatic in comparisons of the DIVF model versus models with static IVF's. For example, the DIVF model generates a pricing error standard deviation of \$.99 compared with about \$2.50 for the three IVF models. The DIVF model generates an implied volatility error standard deviation of 2.34% compared with 7.19%, 6.75%, 5.58% for models 1, 2, and 3.

The ad-hoc models offer performance improvements versus their static counterparts. However, it is important to reiterate that these models are not internally consistent. These models do not provide guidance about how to hedge changes in the implied volatility function, since these changes occur outside of the model specification. As in DFW, these models are tested using the performance of the prior week's estimated IVF in fitting the current week's option prices and implied volatilities.

The best performing ad-hoc model is model 6, which is an implied volatility function defined over moneyness and re-estimated weekly. The pricing error standard deviation for model 6 and the DIVF model are nearly the same. This similarity indicates that an ad-hoc framework is not required to capture

the dynamic behavior of option prices. Time variation in the implied volatility function may be effectively modeled by factoring out a single stochastic variable: at-the-money-implied volatility.

Table 3 also reports statistics suggesting that further extensions of the DIVF framework may be necessary. All of the models have minimum and maximum pricing errors which are large. Using these statistics, the DIVF offers the best pricing accuracy, but the dollar pricing error extremes are about -\$7 and \$8. The DIVF model is unbiased in terms of fitted implied volatilities but is slightly biased in terms of pricing errors, with a mean pricing error of \$.09.

IV. Conclusions

This paper proposes and implements a generalization of implied volatility function models — the dynamic implied volatility function model (DIVF) — which separately models the time-invariant implied volatility function and stochastic state variables. The relative implied volatility function is defined as a time-invariant function of standardized moneyness. Time-variation in the level of implied volatilities is incorporated through the at-the-money implied volatility which scales the relative implied volatility function. At-the-money implied volatility dynamics are modeled as a function of lagged volatility and a non-linear function of the underlying asset return.

Tests based on the full history of S&P500 futures option data show that DIVF outperforms time-invariant IVF models. Performance is somewhat better than an ad-hoc IVF model which is continuously updated; however, a significant advantage of the DIVF model is that it explains option price behavior in an internally consistent manner. This is especially encouraging, considering the parsimonious nature of the DIVF model.

However, it is clear that the dynamics of option prices have not been fully explained by the DIVF model. Generalizations of the DIVF model to incorporate additional state variables such as stochastic relative risk aversion (Rosenberg and Engle, 1999) and more complex at-the-money volatility dynamics offer promising areas for future research.

Bibliography

- Ait-Sahalia, Y. and A. W. Lo, 1998, "Nonparametric Estimation of State-Price Densities Implicit in Financial Asset Prices," *Journal of Finance*, 53, 499-547.
- Barone-Adesi, G. and R. E. Whaley, 1987, "Efficient Analytic Approximation of American Option Values," *Journal of Finance*, 42, 301-320.
- Bates, D., 1995, "Hedging the Smile," Manuscript.
- Black, F., 1976, "The Pricing of Commodity Contracts," *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, 637-654.
- Breeden, D. T. and R. H. Litzenberger, 1978, "Prices of State-Contingent Claims Implicit in Option Prices," *Journal of Business*, 51, 621-651.
- Britten-Jones, M. and A. Neuberger, 1999, "Option Prices, Implied Price Processes, and Stochastic Volatility," Manuscript.
- Dumas, B., J. Fleming and R. E. Whaley, 1998, "Implied Volatility Functions: Empirical Tests," *Journal of Finance*.
- Engle, R. F. and V. K. Ng, 1993, "Measuring and Testing the Impact of News on Volatility," *Journal of Finance*, 48, 1749-1778.
- Glosten, L. R., R. Jagannathan and D. E. Runkle, 1993, "On the Relation Between the Expected Value and the Volatility of the Nominal Excess Return on Stocks," *Journal of Finance*, 48, 1779-1801.
- Jackwerth, J. and M. Rubinstein, 1996, "Recovering Stochastic Processes from Option Prices," Manuscript.
- Natenberg, S. "Option Volatility & Pricing: Advanced Trading Strategies and Techniques." Burr Ridge: Irwin Professional Publishing, 1994.
- Reiner, E., 1999, "Volatility Rules and Implied Processes," Manuscript.
- Rosenberg, J. and R. F. Engle, 1999, "Empirical Pricing Kernels," NYU - Stern School of Business, Manuscript.
- Shimko, D. C., 1993, "Bounds of Probability," *Risk*, 6, 33-37.
- Skiadopoulos, G., S. Hodges and L. Clewlow, 1999, "The Dynamics of the S&P500 Implied Volatility Surface," Manuscript.
- Tompkins, R. "Options Analysis: A State of the Art Guide to Options Pricing, Trading, and Portfolio Applications." Irwin, 1995.

Table 1

Data summary

Characteristics of S&P500 futures options used in the analysis

Variable	N	Mean	Std. dev.	Minimum	Maximum
Time until expiration	53714	54.81	25.00	10.00	100.00
Exercise price	53714	465.85	203.90	145.00	1050.00
Implied standard deviation	53714	365.30	7.32	6.80	89.89
Proportional moneyness	53714	-3.31%	7.95%	-24.99%	25.00%
Standardized moneyness	53714	-0.64	1.44	-5.00	3.80

*There are 3140 dates used in the analysis.

This table presents the S&P500 futures option data used in the paper. Daily data is obtained from the Futures Industry Institute for the period 1983-1998. Implied volatilities (standard deviations in annualized percentage terms) are calculated using the Black (1976) model with the riskless rate proxied by the three month BBA LIBOR rate and the settlement price of the corresponding S&P500 futures contract used as the futures price. The time-until-expiration is measured as the number of calendar days until the expiration date. Proportional moneyness is $K/F_t - 1$, and standardized moneyness is $(K/F_t - 1)/(\sigma_{ATM,t}\sqrt{(T-t)})$.

Table 2
Implied volatility function models

Implied volatility functions - cubic polynomial approximation

	Intercept	a1	a2	a3	Adjusted r2
Model 1: IVF (Black-Scholes)	0.0085				
Model 2: IVF (exercise price)	0.0233	-7.01E-05	1.04E-07	-4.54E-11	11.43%
Model 3: IVF: (prop. moneyness)	0.0082	-0.0181	0.2133	0.5058	39.55%
Model 4: DIVF	-0.1495	0.0471	0.0050	0.0006	91.01%

Bold type denotes significance at the .01 level

Ad hoc implied volatility functions (weekly re-estimation) - average values

	Intercept	a1	a2	a3	Adjusted r2
Model 5: IVF (Black-Scholes)	0.0084				
Model 6: IVF (exercise price)	-0.2090	1.35E-03	-1.83E-06	-1.19E-09	98.20%
Model 7: IVF: (prop. moneyness)	0.0084	-0.0213	0.1288	0.3552	98.31%

DIVF model of at-the-money implied volatility dynamics

	ω	α	β	Adjusted r2
Estimate	0.0006	0.89	0.05	0.9017

Bold type denotes significance at the .01 level

Fitted implied standard deviations (Wednesdays)

	N	Mean	Std	Minimum	Maximum
Model 1: IVF (Black-Scholes)	10855	16.18%	0.00%	16.18%	16.18%
Model 2: IVF (exercise price)	10855	19.08%	2.46%	16.41%	29.04%
Model 3: IVF: (prop. moneyness)	10855	19.10%	4.56%	15.01%	47.08%
Model 4: DIVF	10855	18.98%	6.87%	8.04%	68.38%
Ad-hoc models (weekly re-estimation)					
Model 5: IVF (Black-Scholes)	10855	16.34%	5.75%	8.49%	63.68%
Model 6: IVF (exercise price)	10855	18.72%	7.06%	1.00%	75.48%
Model 7: IVF: (prop. moneyness)	10855	18.83%	7.11%	1.89%	75.75%
Actual	10855	18.96%	7.19%	6.80%	76.41%

This table reports estimates of implied volatility functions for models 1 through 7. For models 1 - 4, daily S&P500 option implied volatilities over the period 1983-1998 are aggregated to estimate a single time-invariant model. Models 5 through 7 are ad-hoc, and estimates are obtained using a weekly fit (with Wednesday data) to the implied volatility functions defined for models 1 - 3. The DIVF at-the-money implied volatility process is estimated using a regression of Wednesday at-the-money implied volatility on its first lag and on the one-week S&P500 return. The implied volatility models are estimated with implied volatility in daily units, which is reflected in the parameter values. All other implied volatilities in the paper are reported in annualized terms by scaling the daily implied volatility by the square-root of 365.

Table 3
Implied volatility function model tests

Implied volatility errors (Wednesdays)

	N	Mean	Std. dev.	Minimum	Maximum
Model 1: IVF (Black-Scholes)	10855	2.78%	7.19%	-9.38%	60.23%
Model 2: IVF (exercise price)	10855	-0.11%	6.75%	-16.89%	50.04%
Model 3: IVF: (prop. moneyness)	10855	-0.14%	5.58%	-21.77%	47.85%
Model 4: DIVF	10855	-0.01%	2.34%	-14.31%	20.92%
Ad-hoc models (weekly re-estimation)					
Model 5: IVF (Black-Scholes)	10855	2.62%	5.19%	-24.35%	43.69%
Model 6: IVF (exercise price)	10855	0.25%	2.51%	-29.32%	43.36%
Model 7: IVF: (prop. moneyness)	10855	0.13%	2.79%	-38.24%	42.64%

Bold-type indicates significantly different from 0 at the .01 level

Dollar pricing errors (Wednesdays)

	N	Mean	Std. dev.	Minimum	Maximum
Model 1: IVF (Black-Scholes)	10855	\$0.37	\$2.67	-\$7.96	\$20.26
Model 2: IVF (exercise price)	10855	-\$0.76	\$2.46	-\$9.54	\$15.53
Model 3: IVF: (prop. moneyness)	10855	-\$0.10	\$2.53	-\$6.94	\$19.92
Model 4: DIVF	10855	\$0.09	\$0.99	-\$6.64	\$8.45
Ad-hoc models (weekly re-estimation)					
Model 5: IVF (Black-Scholes)	10855	\$0.43	\$1.59	-\$11.08	\$14.34
Model 6: IVF (exercise price)	10855	\$0.05	\$1.03	-\$17.64	\$13.97
Model 7: IVF: (prop. moneyness)	10855	\$0.00	\$1.25	-\$10.34	\$14.26

Bold-type indicates significantly different from 0 at the .01 level

This table reports pricing tests of implied volatility functions for models 1 through 7. Each Wednesday a cross-section of option contracts is extracted from the sample and implied standard deviations are calculated. Fitted implied volatilities using the estimated implied volatility functions are obtained using the contemporaneous underlying price, and for the DVF model, using the lagged at-the-money implied volatility and one-week return. Fitted option prices are obtained using the Black (1976) model and the current underlying price, riskless rate, contract characteristics, and fitted implied volatility. Price and standard deviation differences are obtained by subtracting the fitted value from the actual value.

Figure 1
Comparison of actual and fitted at-the-money implied volatility


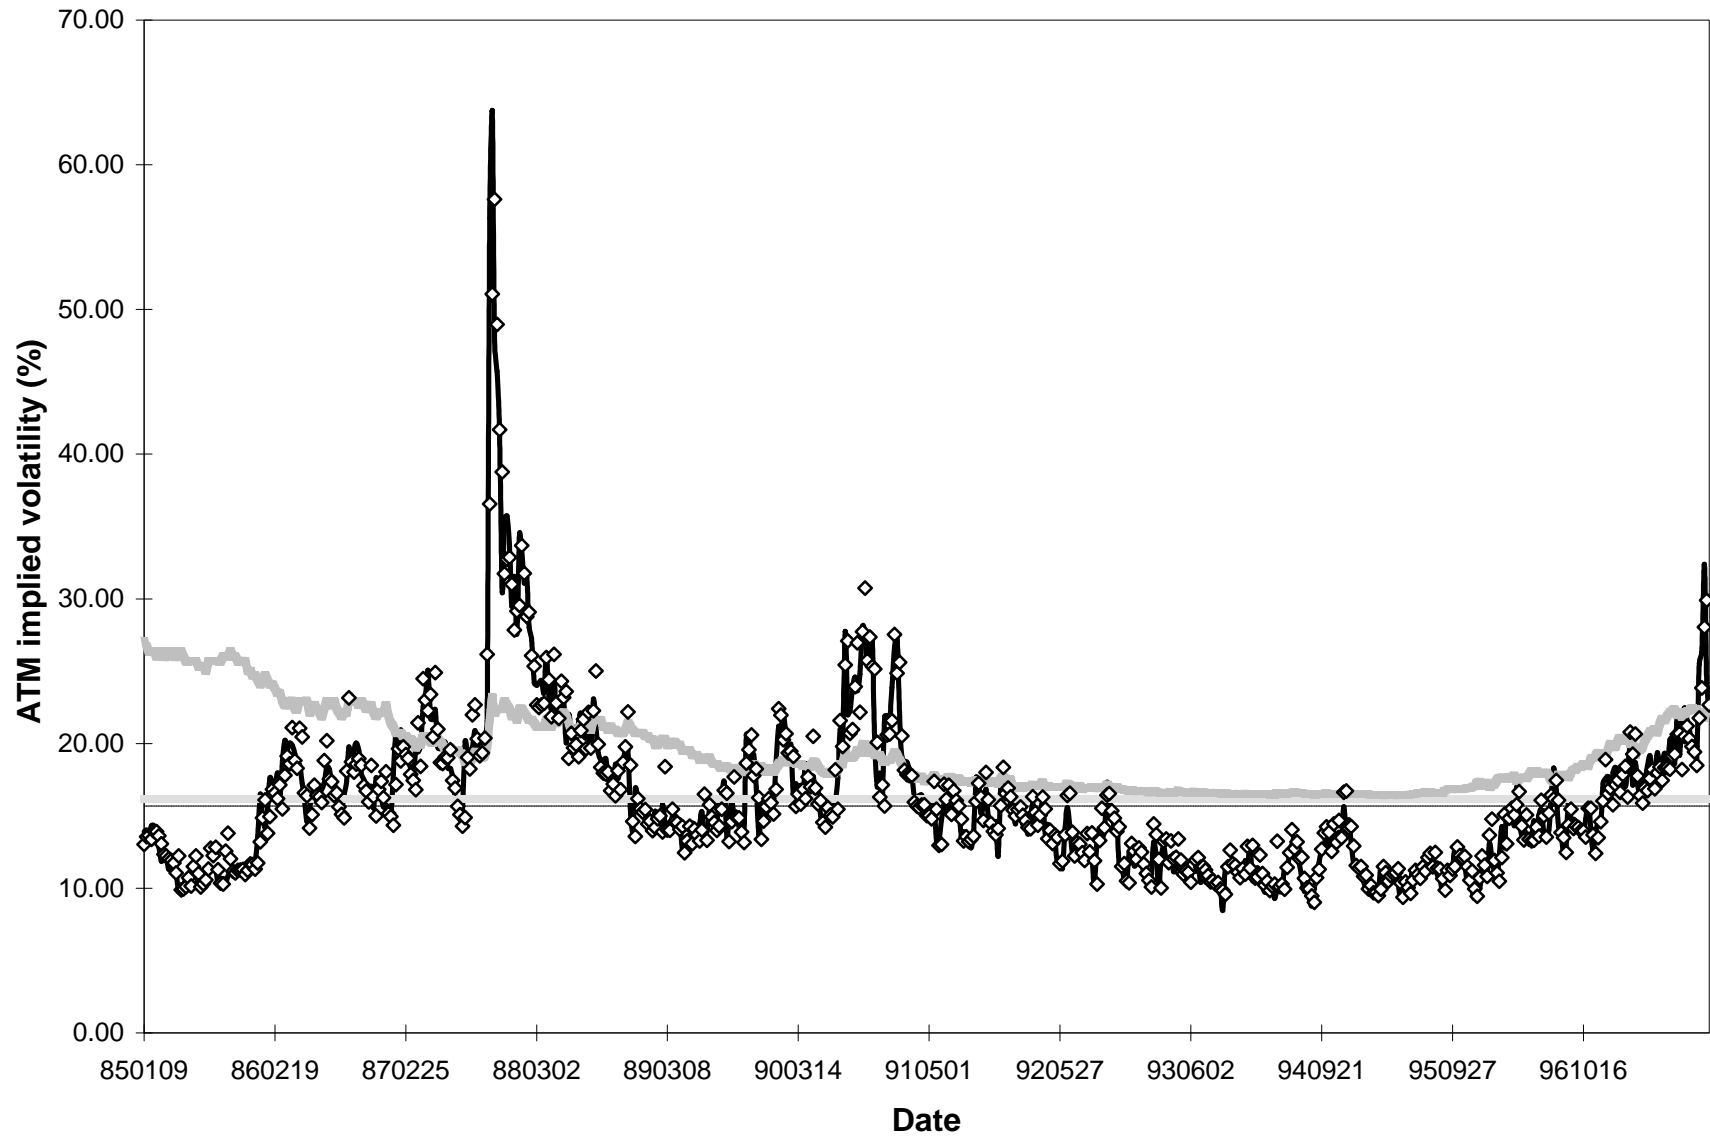


Figure 2
Model 2: Implied volatility function
Implied volatility versus exercise price
S&P500 futures options (1983–1998)




Figure 3
Model 3: Implied volatility function
Implied volatility versus proportional moneyness
S&P500 futures options (1983–1998)




Figure 4
Model 4: Implied volatility function
Relative implied volatility versus standardized moneyness
S&P500 futures options (1983–1998)

