Draft of May 2, 2014
forthcoming in The Review of Finance

Modeling the Dynamics of Correlations among Implied Volatilities

Robert Engle* and Stephen Figlewski**

Implied volatility (IV) reflects both expected empirical volatility and also risk premia. Stochastic variation in either creates unhedged risk in a delta hedged options position. We develop EGARCH / DCC models for the dynamics of volatilities and correlations among daily IVs from options on 28 large cap stocks. The data strong support a general correlation structure and also a 1-factor model with the VIX index as the common factor. Using IVs from stocks that are either highly correlated with the target stock's IV or in the same industry together with the VIX can significantly improve hedging of individual IV changes.

JEL Classifications: G13, G12, C32

Keywords: stochastic volatility, hedging correlation, implied volatility, GARCH

* Professor of Finance
NYU Stern School of Business

** Professor of Finance
NYU Stern School of Business

We thank Amrut Nashikkar and Yinghua Liu for their assistance in the estimations and David Hait of OptionMetrics for providing additional data. Valuable comments and suggestions from Gurdip Bakshi, Liuren Wu, Robert Hodrick, and participants at the 2012 Stern Volatility Institute Conference and the 2012 OptionMetrics conference, as well as two referees and the editor are gratefully acknowledged. We gratefully acknowledge funding support from the NASDAQ OMX Educational Foundation.

Modeling the Dynamics of Correlations among Implied Volatilities

ABSTRACT

Implied volatility (IV) reflects both expected empirical volatility and also risk premia. Stochastic variation in either creates unhedged risk in a delta hedged options position. We develop EGARCH / DCC models for the dynamics of volatilities and correlations among daily IVs from options on 28 large cap stocks. The data strong support a general correlation structure and also a 1-factor model with the VIX index as the common factor. Using IVs from stocks that are either highly correlated with the target stock's IV or in the same industry together with the VIX can significantly improve hedging of individual IV changes.

JEL Classifications: G13, G12, C32

Keywords: stochastic volatility, hedging correlation, implied volatility, GARCH

1. INTRODUCTION

Implied volatility, the volatility value that goes into the market's option pricing formula, is a "risk neutral" value, that impounds both the market's objective estimate of the future realized volatility, and also the market's tolerance for volatility risk. In the Black-Scholes (BS) model, instantaneous volatility is a known constant or, at worst, a nonstochastic function of time. Empirically, stochastic variation in volatility produces risk in a real world delta hedged position in two ways. The underlying stock's realized volatility can be larger or smaller than expected, leading to hedging errors whose size is determined by the option's gamma. Shocks to the market's volatility risk premium also change option prices. Even if an option writer has hedged delta, gamma, and theta, an increase in implied volatility causes an increase in an option's price, and an immediate mark to market loss on the position.

The price sensitivity to a change in implied volatility (IV) is measured by vega, the partial derivative of the (market's) option pricing formula with respect to the volatility parameter. Neither empirical volatility nor risk premia are directly observable and both can change unpredictably over time. However, to the extent that these fluctuations are independent across stocks, vega risk can be reduced through portfolio diversification. The focus in this paper is on modeling the dynamics of implied volatilities across stocks in order to enhance our understanding of vega risk and to develop tools for evaluating and managing it. In the model exposition, "volatility" should be understood as implied rather than realized or empirical volatility.

Figure 1 plots daily time series of implied volatilities for 28 large stocks over the period from Jan. 4, 1996 through Feb. 22, 2009.[1] The stocks are identified by their ticker symbols and the corresponding company names are shown in Table I. Individual stock

[1] These are composite implied volatilities from at the money calls and puts with a weighted average time to maturity of 30 days. The construction method is described more fully below, in Section 4.

IVs vary widely over time, but the fluctuations do not appear to be independent. Distinct patterns are visible, with most stocks showing high values of implied volatility from 2000-2002, low IVs in 2004-06, and very high volatilities at the end of the sample in 2008. This pattern suggests that IVs are partially driven by one or more common factors.

The last graph in the lower right corner of Figure 1 is for the VIX index. The VIX was introduced by the Chicago Board Options Exchange in 1993 as a kind of market implied volatility index. The computation formula for the VIX was revised in 2003; the current VIX reflects an average implied volatility for 30-day options on the Standard and Poor's 500 index.[2] Its time pattern closely resembles those of the individual stocks. We explore below using the VIX index as a proxy for the common factor in explaining the dynamics of IVs and the correlations among them.

At a single point in time, options on a given stock generally exhibit a range of implied volatilities for contracts with a single expiration, that practitioners summarize in the volatility "smile" or "skew." More generally, the "volatility surface" refers to the full set of contracts with all maturities taken together.[3] Attempts to model the dynamics of the volatility surface find evidence of multiple underlying factors related to its level, slope and curvature, as in the term structure of interest rates.[4] However, we will not attempt to incorporate multiple factors into a model of IV dynamics here.

We develop a statistical model which allows changes in implied volatilities, volatilities of implied volatilities and correlations of IVs among a set of underlying stocks, where a stock's IV is measured by a composite of the implied volatilities of near to maturity at-the-money calls and puts. The model correlations can be used to assess the risk of a book of options written on a collection of underlying assets. Such a tool can be useful in developing hedging strategies for a marketmaker's portfolio of delta-gamma hedged positions, as well as for risks associated with more exotic dispersion trades or volatility products such as variance swaps, VIX futures or the newly offered VolContracts^{TM} from VOLX, traded on the CME.

---

[2] See CBOE (2003) for a full description of the VIX index and its method of calculation.

[3] See Gatheral (2006).

[4] See Cont and da Fonseca (2002) or Daglish et al. (2007).

The next section is an overview of the literature. Section 3 develops the GARCH-family models for the dynamics of IVs and their correlations. Section 4 describes the data set, covering 28 stocks in the period 1996-2009, and Section 5 presents the estimation results for the models. Section 6 explores the use of the 1-factor model in designing vega risk hedges for individual stocks by combining positions in the VIX index with IV positions in a small number of related stocks. Section 7 concludes.

2. REVIEW OF THE LITERATURE

The original Black-Scholes (1973) constant volatility framework allows an option to be valued based on a continuously rebalanced delta hedge against the underlying stock. But in the real world large stock price changes expose hedgers to gamma risk that increases with the underlying stock's volatility. Moreover, volatility itself varies stochastically over time, which introduces additional volatility-related risk that is not present, or hedgeable, within the BS model.

Hull and White (1987) and Heston (1993) introduced option models for returns processes with stochastic volatility. Andersen, Benzoni and Lund (2002) and Eraker et al (2003) added stochastic jumps to the returns process and to volatility, and Pan (2002) incorporated them into an option model. Parameter estimation is challenging in these models, however, because volatility is an unobservable stochastic latent factor.

As an alternative model of time-varying volatility, the GARCH framework has a major advantage that there is only a single source of risk.[5] The shocks to returns are also the shocks to returns variance, so there are no latent factors. Both the returns equation and the variance equation are fitted directly from observed returns. The common finding of strong negative correlation between stock returns and returns variances is captured in a GARCH specification by an asymmetry term that increases future volatility more following a negative return than a positive return of the same size, as in the Glosten-Jagannathan-Runkle (1993) version of GARCH, or Nelson's (1991) EGARCH model.

In standard theory, stock risk can be reduced through diversification, but not eliminated because individual stocks covary strongly with a common market factor. By

[5] See Engle (1982), Bollerslev (1986), or Barone-Adesi et al. (2008).

contrast, in Black-Scholes an option can be perfectly hedged by taking a dynamic offsetting position in its underlying stock. Correlation between the underlying stock's return and returns on other stocks, or the market as a whole, only become important for options based on multiple underlying stocks.[6,7] But once volatility is allowed to be stochastically time-varying, delta hedging leaves unhedged volatility risk. Combining delta hedges into a portfolio can mitigate this risk, but it cannot be eliminated if implied volatility shocks are correlated across stocks. In that case, it may even be appropriate to incorporate an option's sensitivity to "nondiversifiable" vega risk into a theoretical pricing framework for options, similar to the Capital Asset Pricing Model for stocks, although we do not attempt to do that here.

GARCH-family models that can apply to a portfolio of stocks must allow the individual pairwise correlations to be time-varying along with the volatilities. The problem is that with n stocks, there are $(n^2 - n)/2$ distinct correlations. To keep the number of fitted parameters from getting out of hand, various constrained models have been developed, with the most popular being the class of models with dynamic conditional correlation (DCC). See, for example, Engle (2002) or Engle (2009).

The discussion and articles cited so far have focused on modeling the underlying stock's "true" or "empirical" or "real world" returns process. Under BS assumptions, risk preferences don't affect the option value directly; the only difference between the empirical distribution and the risk neutralized distribution is a difference in their means. The volatility input to the model should be the market's best estimate of the square root of the stock's integrated return variance from the present to option expiration, which in Black-Scholes is the same under the risk neutral measure. In that case, the implied volatility (IV) obtained from an option's market price by inverting the BS model is the market's objective forecast of future volatility. But it was clear from the beginning that this does not work in the real world. The ubiquitous volatility smile shows that option pricing under the risk neutral measure and under the real world empirical measure are not the same.

---

[6] See Stulz (1982) or Borovkova et al. (2007), for example.

[7] By contrast, derivatives based on credit portfolios depend very heavily on the correlation in default risk among the component securities.

Risk exposures that cannot be perfectly hedged can be expected to be priced in the market. Stochastic volatility is one such risk factor, and as Bakshi and Kapadia (2003), Carr and Wu (2008), and others have shown, implied volatility is on average well above the underlying stock's realized volatility over the option's lifetime. The market price of volatility risk is positive.[8]

Option prices impound the dynamics of returns under the real world measure and also time variation in the market's tolerance for option-related risk. Thus, a rise in expected real world volatility will increase IV and option value. But an option's price can also move without any change in the stock price or the true expected volatility if risk attitudes change. In a delta hedge, any unexpected change in IV produces an option price shock. As Green and Figlewski (1999) amply documented, returns on delta-hedged option positions in the real world are exposed to substantial fluctuations in their supposedly riskless returns due to vega risk.

If delta hedging is not riskless, options market makers and stock portfolio managers are now in a similar situation. Option positions can be delta hedged against stock price changes, but vega risk has to be managed. Dynamic models of the volatility surface have been developed to predict expected changes in IVs, see Gatheral (2006) or Derman et al. (1995) for instance, and unhedgeable vega risk must be diversified. The same issue as with stock portfolios arises with options: Is vega risk "idiosyncratic" so that its effect is largely eliminated in a broadly diversified portfolio of delta hedges, or is there a common "market" IV factor that will always remain?[9] And even without a common factor, as long as IV shocks are not perfectly independent across options, knowledge of

[8] The market price of volatility risk is often said to be negative, not positive. In most cases, this is simply a different way of describing the same phenomenon. Risk neutral volatility is higher than empirical volatility, which causes the prices of options to be higher than if there were no volatility risk premium. Investors pay a positive premium to buy options, because their positive exposure to volatility is a hedge of volatility risk. The returns on those high-priced options will then be lower on average, so the effect of the volatility risk premium on option returns is negative.

[9] Driessen et al. (2009) find that only market-wide volatility risk is priced, while it is not priced at the individual stock level.

their covariance structure allows more efficient diversification through construction of minimum vega risk portfolios of delta hedges.

In 1993, the Chicago Board Options Exchange introduced the VIX index to measure the implied volatility from options on the overall stock market. Given its apparent connection to market risk attitudes, the VIX is often known as the market's "fear gauge."[10] The VIX has been found to be a viable measure of the common component of IV fluctuations for individual options and portfolios of options. Recently introduced futures contracts and options tied to the VIX allow direct hedging of the vega exposure to this common factor.

While some research has been done on modeling the connection between the VIX and IVs from options on individual stocks to support "dispersion" trading, ours is the first paper to apply multivariate GARCH methodology to explore the dynamics of changes in IVs and the correlations among them.[11]

3. THE MODEL

Risk managers, options traders and strategists must understand the risk exposure of an options portfolio that includes contracts written on a number of underlying assets with various strikes and maturities. By invoking a theoretical pricing model, such as Black Scholes, it is easy to evaluate the Greek letter risks of a portfolio of options with a single underlying asset. Traders aim to reduce risk by holding hedged portfolios that are approximately unaffected by small movements in the underlying asset prices and, ideally, in their volatilities. Delta hedging an option position against the underlying stock is relatively easy; vega hedging is not. With multiple underlyings, only the portfolio deltas are typically evaluated.

[10] See Whaley (2000).

[11] See, for example, Jacquier and Slaoui (2007). The CBOE (2009) recently introduced a new Implied Correlation Index based on the dispersion between the VIX and the average implied volatility among options on the stocks that make up the index.

Consider a collection of options whose prices at time t are given by a vector p_t. The price vector for the underlying assets, arranged in the same fashion, is given by s_t and the (Black Scholes) implied volatilities are stacked into a vector v_t. The delta matrix for this set of options is defined by

$$\Delta_t = \frac{\partial p_t}{\partial s_t'} \tag{1}$$

Most elements of this matrix are zero, in theory, since the price of an option on one asset should be unaffected by a change in another asset's price as long as the underlying is unchanged. There may be additional parameters in delta, and each of these must be evaluated at the time the hedge or risk measure is computed. A portfolio with a vector W denoting W_i units of each option i would be valued at $\tilde{\pi}_t = W' p_t$. To make this portfolio delta neutral, offsetting positions would be taken in the underlying assets to give a value for the hedged portfolio of

$$\pi_t = W'\left(p_t - \Delta_t s_t\right) \tag{2}$$

In theory this portfolio has no risk exposure from small movements in any of the components of s_t.

When implied volatility changes, the option prices will also change. The vega matrix for the vector of options is defined as:

$$\Lambda_t = \frac{\partial p_t}{\partial v_t'} \tag{3}$$

If there is only one option for each underlying, Λ_t would be expected to be a diagonal matrix. With multiple options per stock, it would become block diagonal. Since the derivative of s_t with respect to v_t is zero, the vega for the delta neutral portfolio in (2) is

$$\frac{\partial \pi_t}{\partial v_t'} = W'\Lambda_t \tag{4}$$

Only if $W'\Lambda_t = 0$, will this portfolio be independent of the dynamics of the implied volatilities. In theory this can be achieved in a hedge that offsets the volatility exposure with respect to each of the underlyings, but this is rarely possible in practice. If the changes in individual implied volatilities are highly correlated then a short volatility position in one asset can be hedged by a long volatility position in another and the vega risk in the overall portfolio can be minimized. This leads to a need for a dynamic covariance matrix of the implied volatilities of the underlyings. The focus of the paper is on developing expressions for that IV covariance matrix.

The volatility in the returns equation of a typical GARCH model is a forecast of the volatility of returns over the next day. For long horizon forecasts, the expected average volatility of volatility for returns becomes very small as the new information has a relatively small effect on the long run average forecasts. We model the behavior of the volatility of IV over the next day. But implied volatility is a forecast of the average return volatility over the remaining life of an option, which will generally be many days or even years. Even so, we do not necessarily expect volatility of volatility for IVs to become very small at longer horizons, because IVs impound risk preferences in addition to volatility expectations. Changes in market risk premia may have a large effect on long maturity contracts.

A model for the dynamics of implied volatilities will provide a method for calculating the conditional covariance matrix for a set of delta-hedged equity option positions. As part of this model, the covariance matrix will be dynamic, with different correlations in the short and long run. Furthermore, allowing mean reversion in volatilities will produce some predictability in the IVs, although probably not in the option returns themselves. The theoretical models we develop below build up this covariance matrix element by element by specifying the dynamics of each single option's implied volatility, and the correlation between the IVs for a pair of options written on different stocks.

For underlying asset i, the one month at the money implied volatility on day t is denoted $v_{i,t}$. We consider the statistical model for the change in IV

$$\log(\frac{v_{i,t}}{v_{i,t-1}}) = \psi_{i,t} + \sigma_{i,t}\varepsilon_{i,t}, \quad \varepsilon_{i,t} \sim IID(0,1) \tag{5}$$

where this notation indicates that the errors are independent over time with mean zero and variance one. The functions $(\psi_{i,t}, \sigma_{i,t})$ are the conditional mean and standard deviation of the change in log IV and are measurable with respect to information at t-1. Note that we are now changing from the continuous-time setting of theoretical option pricing models to the discrete-time framework of volatility modeling under GARCH.

We model volatility as a predictable mean reverting process. A simple specification for the expected change in log volatility is an AR(1):

$$E_{t-1} \log\left(\frac{v_{i,t}}{v_{i,t-1}}\right) = \theta_{1,i} + \theta_{2,i} \log(v_{i,t-1}) \tag{6}$$

where the subscripts indicate that the theta parameters can vary over names but not over time.

3.1 VOLATILITY OF VOLATILITY

The conditional variance of the log of the implied volatility, $\sigma^2_{i,t}$, is a measure of the volatility of (implied) volatility. Its dynamics can be parameterized naturally in the same way as is done for the return volatilities of regular assets. One difference is that, unlike security returns, the change in implied volatility can have a non-trivial mean term that must be taken out before the volatility of volatility is measured. In this study we use the EGARCH volatility process, as it easily accommodates the non-standard asymmetry in volatility that is found in virtually all volatility of volatility models.

$$\log\left(\sigma^2_{i,t}\right) = \phi_{1,i} + \phi_{2,i}\left|\varepsilon_{i,t-1}\right| + \phi_{3,i}\varepsilon_{i,t-1} + \phi_{4,i}\log\left(\sigma^2_{i,t-1}\right) \tag{7}$$

Standard TARCH specifications which differ from (7) by having the variance on the left hand side rather than its log, often maximize the likelihood function with a negative value

of the parameter $\phi_{3,i}$ that measures the asymmetric part of volatility. In such a model, there is a possibility of a negative volatility, especially during the estimation procedure. That is not a problem with EGARCH.

It is very common with financial assets to find that negative returns have a bigger impact on future volatility than positive returns. However for options, declines in IV are found to have a smaller impact on future volatility of volatility than increases do. This is consistent with the familiar positive skew in returns on volatility contracts, which is opposite to the skew observed in most asset classes.[12] Consequently, we expect $\phi_{3,i} > 0$.

With this volatility process it is simple to model jointly the mean and the volatility of each individual log implied volatility. The final specification decision is what distribution to assume for the $\varepsilon_{i,t}$ residuals. The normal is a standard choice, but it is found that financial data are more consistent with fatter-tailed distributions. In the empirical results presented below, we have assumed a Student-t distribution, with the degrees of freedom, which determines tail fatness, as a parameter to be fitted for our stocks. Typical degrees of freedom values turn out to be around 4, which implies quite fat tails.

3.2 CORRELATIONS

The conditional correlations between two implied log volatilities are given by

$$\rho_{ij,t} = \frac{E_{t-1}\left(\log\left(v_{i,t}/v_{i,t-1}\right)-\psi_{i,t}\right)\left(\log\left(v_{j,t}/v_{j,t-1}\right)-\psi_{j,t}\right)}{\sqrt{Var_{t-1}\left(\log\left(v_{i,t}/v_{i,t-1}\right)\right)Var_{t-1}\left(\log\left(v_{j,t}/v_{j,t-1}\right)\right)}} \tag{8}$$

Direct substitution from (5) implies that

$$\rho_{ij,t} = E_{t-1}\left(\varepsilon_{i,t}\varepsilon_{j,t}\right) \tag{9}$$

[12] Also, as pointed out in Engle (2011) and Engle and Mistry (2007), this is likely to be a property of any hedging portfolio since increases in the level of volatility are likely to signal that investors need a larger hedge and hence need to purchase additional volatility exposure, thereby driving up its price.

Interestingly, the conditional correlations between log IVs are simply the correlations between the disturbances. These correlations will be estimated with dynamic structure so that they can evolve over time.

The Dynamic Conditional Correlation or DCC model of Engle (2002, 2009) is ideally suited to this problem. In that model, the dependent variables are first adjusted for means and volatilities and then the correlations are computed for the standardized returns, which are interpreted as the epsilons. The DCC model estimates a full set of correlations at one time with a very restricted parameterization.

Letting ε_t be an n x 1 vector of standardized log volatility residuals, the objective is to estimate the correlation matrix which is defined as

$$R_t = E_{t-1}\left(\varepsilon_t \varepsilon_t '\right) \tag{10}$$

The model for these correlations is given by the following three equations:

$$\bar{R} = \frac{1}{T}\sum_{t=1}^{T} \varepsilon_t \varepsilon_t ' \tag{11}$$

$$Q_t = \bar{R} + \alpha\left(\varepsilon_{t-1}\varepsilon_{t-1}' - \bar{R}\right) + \beta\left(Q_{t-1} - \bar{R}\right) \tag{12}$$

$$R_t = diag\left(Q_t\right)^{-1/2} Q_t diag\left(Q_t\right)^{-1/2} \tag{13}$$

Equation (11) calculates the unconditional correlations of the ε_t's. Equation (12) creates a matrix Q_t that evolves based on just two unknown parameters α and β which govern the decay of correlation information. The Q_t matrix is approximately a correlation matrix but it does not have ones on the diagonal every time period. Equation (13) converts the Q_t matrix into a proper correlation matrix R_t and is the only place where a non-linearity is introduced. This non-linearity is a natural consequence of estimating correlations which are bounded as a function of data which are not.

The two unknown parameters of the DCC model can be estimated with maximum likelihood. However, substantial research shows that there are biases for large systems

that let alpha approach zero. Several alternative estimation strategies have been proposed and in this analysis one of the simplest but very effective strategies called the MacGyver method is applied, see Engle (2009, 2009a). This method estimates n(n-1)/2 bivariate models, as in Equation (14).

$$\bar{R}_{i,j} = \frac{1}{T}\sum_{t=1}^{T} \varepsilon_{i,t}\varepsilon_{j,t} \tag{14}$$

$$Q_{ij,t} = \bar{R}_{i,j} + \alpha_{i,j}\left(\varepsilon_{i,t-1}\varepsilon_{j,t-1} - \bar{R}_{i,j}\right) + \beta_{i,j}\left(Q_{ij,t-1} - \bar{R}_{i,j}\right) \tag{15}$$

$$R_{ij,t} = \left(Q_{ii,t}\right)^{-1/2} Q_{ij,t} \left(Q_{jj,t}\right)^{-1/2} \tag{16}$$

without imposing the restriction that alpha and beta must be the same for all i and j. Then the medians of these two series are computed and used in (15)-(16) in place of $\alpha_{i,j}$ and $\beta_{i,j}$ to compute all the correlations. In this way outliers are ignored and usable estimates are obtained even if some of the bivariate pairs fail to converge or lie outside the natural (0,1) interval. This method is compared with composite likelihood by Engle et al. (2008).

3.3 A ONE-FACTOR MODEL

Empirically it is clear that the correlations between implied volatility innovations are much smaller than the (unconditional) correlations between IV levels. For some periods, all IVs are above the long run average and at other times none are. This suggests a factor model or some exogenous driving variables. The models for the $\psi_{i,t}$ for each asset would depend upon variables from the other assets, so that short run correlations can become bigger due to the feedback.

To incorporate this effect, we introduce a common term into the mean IV change for each asset. In the same way that the S&P 500 index serves as a proxy for the "market portfolio," a natural choice for a measure of aggregate market implied volatility is the VIX. This index is closely related to the implied volatilities of 30-day options on the S&P500 index and it impounds a forecast of future volatility.

Suppose the changes in log VIX satisfy the same model discussed above.

$$\log\left(\frac{VIX_t}{VIX_{t-1}}\right) = \psi_{VIX,t} + \sigma_{VIX,t}\varepsilon_{VIX,t}$$

$$= \theta_{1,VIX} + \theta_{2,VIX}\log(VIX_{t-1}) + \sigma_{VIX,t}\varepsilon_{VIX,t} \quad (17)$$

$$\log(\sigma^2_{VIX,t}) = \phi_{1,VIX} + \phi_{2,VIX}\left|\varepsilon_{VIX,t-1}\right| + \phi_{3,VIX}\varepsilon_{VIX,t-1} + \phi_{4,VIX}\log(\sigma^2_{VIX,t-1}) \quad (18)$$

The research on stochastic volatility provides strong evidence that returns variances are mean-reverting. In (17), the long term mean will be absorbed in the constant term $\theta_{1,VIX}$, , and $\theta_{2,VIX}$ should be negative.

Consider now the following model for the dynamics of a single stock's v_i, where mean reversion for the individual stock's log variance and the VIX and are incorporated in the $\theta_{3,i}$ and $\theta_{4,i}$ terms:

$$\log\left(\frac{v_{i,t}}{v_{i,t-1}}\right) = \theta_{1,i} + \theta_{2,i}\log\left(\frac{VIX_t}{VIX_{t-1}}\right) + \theta_{3,i}\log(v_{i,t-1}) + \theta_{4,i}\log(VIX_{t-1}) + \sigma_{i,t}\,\varepsilon_{i,t} \quad (19)$$

In this specification, the variance equation is still given by Equation (7), but $\varepsilon_{i,t}$ now refers only to the volatility of the change in v_i that is independent of the change in VIX.

This 1-factor model permits economy-wide information as reflected in the VIX to influence volatilities in all names simultaneously. The contemporaneous relation between VIX and each name provides an explanation for the correlation in IVs. Thus the second coefficient $\theta_{2,i}$ is the volatility equivalent of the beta that is routinely calculated in analyses of equity returns. In addition, both v_i and VIX are mean-reverting, at possibly different rates. In this specification, the constant term absorbs the long run means while the terms multiplied by $\theta_{3,i}$ and $\theta_{4,i}$ capture the effects of reversion from the date t-1 levels toward those long term means. Given the way $v_{i,t-1}$ and VIX_{t-1} enter Equation (19), we expect $\theta_{3,i} < 0$ and $\theta_{4,i} > 0$. To the extent that an individual stock's IV and the VIX are cointegrated, mean reversion here would essentially be reversion toward their long term relationship. In that case, one might expect these coefficients also to be of similar magnitude. We find both of these properties in the coefficient estimates reported in Table III below.

4. DATA

We will proceed to estimate the 0-factor model described above in Equations (5)-(7), and the 1-factor model of (17)-(19) using daily returns for a sample of large stocks. The time period covered is January 4, 1996 through February 27, 2009. To create a sample of large actively traded stocks, the firms originally selected for the study were the 30 members of the Dow Jones Industrial Average. Two stocks were eliminated due to mergers and other changes in index composition, leaving 28 firms, with a total of 378 pairwise correlations among them. The list of firms and their ticker symbols is shown in Table I. The implied volatilities used in the study are drawn from the OptionMetrics database, available through the Wharton Research Data Services (WRDS) system.

The GARCH models for individual stock volatilities are fitted to standardized IVs for at the money options with maturities of 30 calendar days. The methodology is essentially the same as for the original computation of the CBOE's VIX index, using weighted averages of IVs from traded options. Specifically, to get the standardized call IV, OptionMetrics averages the IVs from the two calls with the least out of the money and the least in the money strikes, for the longest available maturity less than 30 days. These IVs are weighted so that the weighted average of the strikes equals the current stock price. A similar averaging is done for calls with the same strikes and the shortest maturity greater than or equal to 30 days. The two blended IVs are then averaged using weights such that the weighted average maturity is 30 days. This gives a composite implied volatility for call options. The equivalent calculation is then done with four put options whose strikes bracket the current forward price for the stock and whose maturities bracket 30 days, to get a put IV. We then take a simple average of these OptionMetrics standardized call and put IVs to get a single overall IV per stock per day which becomes the $v_{i,t}$ series in the model.

In the 1-factor model, the dynamics of individual stock IVs are assumed to be correlated with the actual (new) VIX index reported by the CBOE as a proxy for the

aggregate volatility of the whole stock market.[13] Data on the VIX index was downloaded from the Chicago Board Options Exchange website.

5. EMPIRICAL ESTIMATION OF THE CORRELATION MODELS

Each element of the correlation matrix R_t defined in (14)-(16) corresponds to a pair of options, whose IVs are assumed to vary stochastically over time. The individual stock volatilities are assumed to follow either (5)-(6), in the 0-factor model, or (19), in the 1-factor model. The estimated parameters for the two models are presented in Tables II and III.

In Table II, the first two lines show the averages across all stocks of the individual coefficients and their t-statistics. The next two lines fit the 0-factor model for the VIX index, followed by results for the 28 individual stocks. All coefficients are highly statistically significant on average. The mean value of the autoregressive coefficient θ_2 is significantly negative but very small, indicating only a mild degree of mean reversion in implied volatilities at a daily interval, for both the average stock in the sample and for the VIX index. Among the individual stocks, 26 out of 28 of the θ_{2i} estimates are significantly negative, and only 2 are positive: 0.002 for AXP (American Express) and 0.001 for BAC (Bank of America). We suspect that this result for the two large banks may be an artifact resulting from the instability in the banking sector in late 2008 and 2009.

The results for the volatility of volatility equation show the expected positive signs on ϕ_2, ϕ_3 , and ϕ_4. Log variance today is strongly related to yesterday's log variance and the size of yesterday's IV shock. The positive and highly significant estimate for the ϕ_3 coefficient, averaging 0.070, indicates asymmetry in the IV process, such that a positive shock increases future IV by more than a negative shock of the same size.

[13] The procedure for computing the VIX index was changed from the original method in 2003. In this study we use the new VIX that is extracted by a nonparametric procedure from out of the money calls and puts over a wide range of strike prices.

Interestingly, the asymmetry coefficient for the VIX is almost twice as large as the average across single stocks. Among the individual equations, every one shows positive estimates for the coefficient on the absolute value of yesterday's shock ϕ_2, none with a t-statistic below 10. Similarly, the coefficient on lagged volatility of volatility, ϕ_4 , is highly significant, with the lowest t-statistic being 49.45. In all but two cases ϕ_3 is positive and it is significant in all but one of these.

Table III presents the estimation results for the 1-factor model, with the (new) VIX index as the proxy for the common factor. As in Table II, every model coefficient was highly significant on average. The average value for θ_2, the sensitivity to the common factor, was 0.325 with an average t-statistic of more than 31. This coefficient was highly significant for every individual stock, ranging in value between 0.205 and 0.526 with a t-statistic above 20 in every case. This provides strong evidence that there is a common factor driving individual stock implied volatilities that can be well proxied by the VIX index.

The estimates for θ_3 were also highly significant and all were negative, indicating mean reversion of the individual IVs toward its long run value, but the small values for the coefficient, the largest in absolute value being -0.040, indicate that the speed of convergence is fairly slow. The θ_4 coefficients on the VIX mean correction term were of comparable size to those for θ_3 but all were positive. A notable difference between the variance equations for the 0-factor and the 1-factor specifications is that both the ϕ_3 and ϕ_4 coefficients are smaller in Table III than in Table II. Idiosyncratic IV variance is less persistent than total variance for individual names and less asymmetrical.

In the EGARCH specification, stability of the system is ensured if ϕ_4 is below 1.0. In both Table II and Table III, estimates of ϕ_4 are less than 1.0 with a high degree of statistical significance in every case. We note also that the t-distributed shocks were found to have quite fat tails, with the degrees of freedom parameter averaging under 4 for both specifications.

The Macgyver procedure computes values for α_{ij} and β_{ij} in (15) across all pairs (i,j), takes the median α and β, and applies them to construct estimates for the dynamics of the pairwise correlations between ε_{it} and ε_{jt} . The resulting time series of the mean correlation across pairs for the 0-factor model and the 1-factor model are plotted in

Figure 2. The variability in average correlation over time is striking. Because the 1-factor model removes the common factor associated with the VIX index, the average correlation for the residuals is both lower and less variable over time than correlations in the 0-factor model that treats the VIX factor as being part of the ε_{it} and ε_{jt} residuals.

6. HEDGING DYNAMIC VEGA RISK

This section illustrates how our analysis can be applied to hedging vega, a derivative security's exposure to the risk of an adverse change in its implied volatility. Existing VIX futures provide linear payoffs proportional to the change in the VIX index, which we have taken as a proxy for the common implied volatility factor. For convenience, in our examples we will assume that there exist comparable linear payoff contracts on individual stock IVs, such that the change in value for contract i is proportional to the change in v_i.[14] This is meant to be just an illustration of how our analysis could be used to set up a vega hedge using a single "market" IV contract and a small number of contracts on individual stock IVs. We do not try to optimize its performance out of sample, for example, or moderate the effect of extreme parameter estimates for some days.

Let P denote a single asset, a portfolio of securities, or any other position whose returns R_P are correlated with those of a set H of hedging securities, R_{Hi}.[15] A standard result from the hedging literature is that the vector of risk minimizing quantities for the hedging securities, W_H^*, is given by

[14] If, instead, the hedging securities are individual stock options or delta-hedges, the following equations must be modified to include the option vegas, in order to convert the changes in the IV's into dollar values.

[15] In this section, we are abstracting from hedging possible exposure to other sources of risk, notably from changes in the prices of the underlying assets. In hedging an actual portfolio of equity options, delta will typically be neutralized and other Greek letter risks may be hedged also. The hedging strategy described here should be used to hedge the exposure to IV changes that are independent of those other risk factors.

$$W_H^* \;=\; Var(R_H)^{-1} Cov(R_P, R_H) \qquad (20)$$

where $Var(R_H)$ is the variance-covariance matrix of the hedging securities' returns and $Cov(R_P,R_H)$ is the vector of covariances between the target portfolio and the hedge security returns.

6.1 HEDGING COMMON FACTOR AND IDIOSYNCRATIC VEGA RISK FOR A SINGLE NAME

In applying our 1-factor correlation model to the vega hedging problem, we will approximate the change in log IV by the percentage IV change:

$$d\log(x) \;=\; \frac{dx}{x} \;\cong\; \frac{\Delta x}{x} \qquad (21)$$

The first equality holds in continuous-time; the second holds approximately over a short discrete time interval Δt. Since our observation intervals are mostly 1 day and never more than 5 calendar days, the approximation is very close.

Equation (19) shows the model dynamics for an individual stock's implied volatility v_i in terms of the current and previous period values of the VIX. Using (21), the stochastic part of Δv_i is

$$\Delta v_{i,t} \;-\; E[\Delta v_{i,t}] \;=\; v_{i,t}\left(\theta_{2,i}\,\sigma_{VIX,t}\,\varepsilon_{VIX,t} \;+\; \sigma_{i,t}\varepsilon_{i,t}\right) \qquad (22)$$

In our single-stock hedging examples, the dimension of the target portfolio of linear payoffs to be hedged is 1, but for greater generality, we show the hedging solution for a portfolio P that may contain N_P different IV exposures. Let H denote the portfolio of N_H hedging assets, and W, v, θ_2, σ, and ε all denote vectors of the corresponding individual parameter values from a set of securities indicated by subscript P or H. For notational convenience we will suppress the dependence of the risk parameters on t, except where it is needed.

The symbol ∘ will denote the Hadamard product, meaning element by element multiplication of two vectors. For example, $\sigma \circ \varepsilon \equiv \text{diag}(\sigma)\,\varepsilon$ produces a vector with individual elements $\sigma_i \varepsilon_i$. Thus, Σ_{PH}, the N_P x N_H covariance matrix for the idiosyncratic portions of the IV changes between the securities in P and those in H is

$$\Sigma_{PH} \equiv E[(v_P \circ \sigma_P \circ \varepsilon_P)(v_H \circ \sigma_H \circ \varepsilon_H)'] \tag{23}$$

The variance-covariance matrices for the target and hedge portfolios, Σ_P and Σ_H, are defined analogously.

The change in portfolio P is given by $W_P' \Delta v_P$ and the change in the hedging portfolio H is $W_H' \Delta v_H$. Both the target portfolio and the hedging portfolio are exposed to common factor risk. This hedging problem differs from the general case because there is a contract on the VIX that has only common factor exposure and no idiosyncratic vega risk. This allows the minimum variance hedge to use derivatives on single stocks to hedge the exposure to idiosyncratic IV changes, and the VIX contract to eliminate the combined common factor risk from the target and the hedging portfolios together.

From (22), the time-varying sensitivity of security i to the date t realization of the common factor is given by $v_{i,t}\, \theta_{2,i}\, \varepsilon_{VIX,t}$. For any values of W_P and W_H, total sensitivity to $\varepsilon_{VIX,t}$ of the overall hedged position is then

$$W_P' \left(v_P \circ \theta_{2,P}\right) + W_H' \left(v_H \circ \theta_{2,H}\right)$$

which can be completely offset by shorting the following amount of the VIX contract

$$W_{VIX}^* = \left(W_P' \left(v_P \circ \theta_{2,P}\right) + W_H^{*\prime} \left(v_H \circ \theta_{2,H}\right)\right)/VIX \tag{24}$$

Using(23) in Equation (20) to solve for the quantities of the hedging assets to sell against portfolio P to minimize the idiosyncratic portion of the risk, W_H^*, gives

$$W_H^* \;=\; \Sigma_H^{-1}\,\Sigma_{PH}\,'\,W_P \qquad (25)$$

The minimum vega risk hedge will have no common factor risk and idiosyncratic vega risk with total variance

$$\sigma^2_{min\,var} \;=\; W_p'\,\Sigma_P\,W_P \;+\; W_H^{*\prime}\,\Sigma_H\,W_H^* \;+\; 2\,W_p'\Sigma_{PH}\,W_H^* \qquad (26)$$

Table IV reports on the first of two experiments to illustrate the performance of hedges constructed using (24) and (25) with the single-factor model parameter estimates for our sample of stocks.

First, in every period t we consider hedging each individual stock i's IV over the next period by constructing a hedge from the VIX contract and implied volatility contracts on the stock with the highest estimated period t idiosyncratic IV correlation with stock i. The table shows the average hedge performance across stocks and over time in terms of the standard deviation in volatility points of the unhedged position, a hedge using only the VIX contract, and an extended hedge with both the VIX and single stock IV contracts.

On average, the one-day standard deviation of the change in IV across our 28 stocks was 0.0227, (2.27 percentage points in annualized volatility). Optimally hedging this vega risk using just the VIX contract would have reduced the average standard deviation of the hedged positions to 0.01971. Adding a position in the most correlated single stock contract would have reduced the hedge standard deviation a little further, to 0.01970.

The statistical significance of this improvement in hedge performance was gauged in two ways. Differences in the squared hedging errors from the different strategies were computed for each single stock, and t-statistics on the mean differences among them were calculated. The third column of the table report the averages of those t-statistics across stocks. The average t-statistic on the difference between the unhedged position and the extended hedge averaged more than 3.0.

Not all stocks showed significant improvement. For 20 of 28 stocks, the extended hedge significantly reduced the standard deviation of the change in its IV, but in 5 cases it increased and the other 3 showed an improvement that was not statistically significant.

The fourth column reports on the difference in squared hedging errors between the extended hedges and hedges using only the VIX index to offset risk exposure to the common IV factor.

The VIX by itself was quite effective, achieving almost as good performance as the extended hedges and with considerably lower variability across names, such that the t-statistic on the improvement over the unhedged position was 7.301. The performance of the extended hedge relative to the VIX only hedge was positive but not significant on average. Even so, the extended hedges were significantly better than VIX only hedges for 10 of the stocks, while in no case did the VIX only hedge significantly outperform the extended hedge.

In order to judge statistical significance for the whole sample together, we created a pooled sample, estimated the average difference in mean squared error, and computed t-statistics using robust standard errors corrected for heteroskedasticity across names and for contemporaneous correlation in the errors.[16]

These results, shown in columns 5 and 6, confirm what the averages across names showed: both the extended hedge and VIX only strategies significantly reduced risk, but extending the hedge by adding volatility contracts on the single name whose correlation was predicted to be the greatest on each date gave rather limited improvement. One reason for this is that simply picking the stock with the highest estimated correlation on each date will tend to select outliers and stocks with positive estimation errors. One way to reduce this problem in practice could be to focus on correlations over longer time windows, for example selecting not the single stock with the highest estimated correlation on date t, but rather the one that was, say, in the top 5 most frequently over the past month.

A more natural way to use the dynamic IV correlation model in hedging is to select hedging stocks that are expected on a priori grounds to be closely correlated with the target portfolio, for example, stocks from the same industry. Our sample is limited to only 28 stocks, but three of them are drug companies, three are banks, and there are three other sets of two stocks each from the same SIC 4-digit industry. Table V explores hedging each of the twelve stocks for which this can be done using the VIX and one or two other stocks from the same industry.

[16] The procedure is implemented in EViews, using White's robust procedure for cross-section estimation.

The improvement in hedging performance is substantial, especially for the hedges with 2 stocks in the same industry. Of the twelve, nine showed significant improvement and only one of the hedges failed. For seven of them, the extended hedge did significantly better than hedging with the VIX alone, and in no case did the VIX only hedge beat the extended hedge significantly.

7. CONCLUDING COMMENTS

Although the original Black-Scholes model did not envision changes in the volatility parameter, stochastic changes in implied volatilities are an important source of risk for real world options marketmakers. Correlation in IV changes prevents this vega risk from being eliminated by diversification in a portfolio of delta hedged options positions.

Empirical volatilities and correlations have been found to be systematically time-varying in ways that can be captured with GARCH-family models. This paper has explored whether correlations in implied volatility changes also show dynamics that can be captured with multivariate GARCH models. We developed two such models using EGARCH, which is better suited to modeling volatilities that must remain nonnegative, and Dynamic Conditional Correlation (DCC) for modeling correlation dynamics. One specification has time-varying correlations driven by shocks that are possibly correlated across stocks but without a single common factor, and in the second model, IV changes are also correlated through exposure to a single market volatility process that we proxied by the VIX index.

Our estimates strongly supported both models in US equity data. There was a slight but significant degree of mean reversion in the "return" equation ($\theta_2 = -0.01$ on average in the 0-factor model). In the 1-factor model, among 28 large stocks, every one showed a strong correlation between its IV change and the change in the VIX, with an average of 0.325 for θ_2 in that model. In the volatility of volatility equations, both specifications showed a strong positive asymmetry. In the 0-factor model, for 25 out of 28 stocks, a positive shock to implied volatility increases future IV volatility significantly

more than a negative shock of the same size. For 19 stocks, the coefficient on lagged variance exceeded 0.9 for the 0-factor model. The idiosyncratic IV changes in the 1-factor model showed less asymmetry and less persistence. The use of the Student t-distribution rather than the normal to model the shocks helped to reduce the impact of outliers. The estimated degrees of freedom averaged under 4.0 across stocks, meaning the distribution had very fat tails.

We described in general how to use our model to construct a dynamic hedge for the vega risk of an option position, or portfolio, with a hedging portfolio of IV exposures. Using the fitted model coefficients to hedge a single option's vega risk with the VIX and volatility contracts on one other name suggested that significant risk reduction could be achieved by offsetting exposure to a common implied volatility risk factor and also hedging the idiosyncratic residuals from the IV dynamics equations. Turning our essentially statistical results into an operational hedging strategy should involve a careful effort to limit the effects of sampling noise and to incorporate non-statistical information such as the target firm's industry into the hedge design.

References

Andersen, T., Benzoni, L., and Lund, J. (2002). An empirical investigation of continuous-time equity return models, *Journal of Finance* **57**, 1239-1284.

Andersen, T., Bollerslev, T., Diebold, F., and Labys, P. (2003). Modeling and forecasting realized volatility. *Econometrica* **71**, 579–625.

Bakshi, G., Cao, C. and Chen, Z. (1997). Empirical performance of alternative option pricing models, *Journal of Finance* **52**, 2003–2049.

Bakshi, G. and Kapadia, N. (2003). Volatility risk premiums embedded in individual equity options: some new insights, *Journal of Derivatives* **11**, 45-54.

Barone-Adesi, G., Engle, R. F., and Mancini, L. (2008). A GARCH option pricing model with filtered historical simulation, *Review of Financial Studies* **21**, 1223-1258.

Black, F. and Scholes M. (1973). The pricing of options and corporate liabilities, *Journal of Political Economy* **81**, 637-654.

Bollerslev, T. (1986). Generalized autoregressive conditional heteroskedasticity, *Journal of Econometrics* **31**, 307–27.

Borovkova, S., Permana F. and van der Weide, H. (2007). A closed form approach to valuation and hedging of basket and spread option, *Journal of Derivatives* **14**, 8-24.

Carr, P. and Wu L. (2009). Variance risk premiums, *Review of Financial Studies* **22**, 1311-41.

Chicago Board Options Exchange (2003). *VIX CBOE Volatility Index,* http://www.cboe.com/micro/vix/vixwhite.pdf.

Chicago Board Options Exchange (2009). *CBOE S&P 500 Implied Correlation Index*, http://www.cboe.com/micro/impliedcorrelation/ImpliedCorrelationIndicator.pdf.

Christoffersen, P., Jacobs, K., and Chang, B. Y. (2013). Forecasting with option implied information, in G. Elliott and A. Timmermann (eds.), *Handbook of Economic Forecasting, Vol. 2*, Elsevier, Amsterdam, Netherlands.

Cont, R., and da Fonseca, J. (2002). Dynamics of implied volatility surfaces, *Quantitative Finance* **2**, 45-60.

Daglish, T., Hull, J., and Suo, W. (2007). Volatility surfaces: theory, rules of thumb, and empirical evidence. *Quantitative Finance* **7**, 507-524.

Demeterfi, K, Derman, E. , Kamal, M., and Zou, J. (1999). A guide to volatility and variance swaps, *Journal of Derivatives* **6,** 9 - 32.

Derman, E., Kani, I., and Zou, J. (1995). *The Local Volatility Surface: Unlocking the Information in Index Option Prices*, Goldman Sachs Research Notes.

Driessen, J., Maenhout, P., and Vilkov, G. (2009). The price of correlation risk: evidence from equity options, *Journal of Finance* **64**, 1377-1406.

Engle, R. (1982). Autoregressive Conditional Heteroscedasticity with estimates of the variance of United Kingdom inflation, *Econometrica* **50**, 987–1007.

Engle, R. (2002). Dynamic conditional correlation: a simple class of multivariate generalized autoregressive conditional heteroskedasticity models, *Journal of Business and Economic Statistics* **20**, 339-350.

Engle, R. (2009). *Anticipating Correlations*, Princeton University Press, Princeton New Jersey.

Engle, R. (2009a). High dimension dynamic correlations, in: J. L. Castle and N. Shephard (eds.), *The Methodology and Practice of Econometrics: A Festschrift in Honour of David F. Hendry*, Oxford University Press, Oxford, U.K.

Engle, R. (2011). Long-term skewness and systemic risk, *Journal of Financial Econometrics*, **9**, 437-468

Engle, R. and Mistry, A. (2014). Priced risk and asymmetric volatility in the cross-section of skewness, *Journal of Econometrics*, forthcoming.

Engle, R., and Rangel, J.G. (2008). The Spline-GARCH model for low-frequency volatility and its global macroeconomic causes, *Review of Financial Studies* **21**, 1187-1222.

Engle, R., Shephard, N., and Sheppard, K. (2008). Fitting and testing vast dimensional time-varying covariance models, unpublished working paper, University of Oxford.

Eraker, B., Johannes, M., and Polson, N. (2003). The impact of jumps in volatility and returns, *Journal of Finance* **58**, 1269-1300.

Gatheral, J. (2006). *The Volatility Surface: A Practitioner's Guide*, Wiley Finance, Hoboken, New Jersey.

Giamouridis, D. and Skiadopoulos, G. (2012). The informational content of financial options for quantitative asset management, in: B. Scherer and K. Winston (eds.), *The Oxford Handbook of Quantitative Asset Management*, Oxford University Press, Oxford, U.K.

Glosten, L., Jagannathan, R., and Runkle, D. (1993). On the relation between the expected value and the volatility of the nominal excess return on stocks. *Journal of Finance* **48**, 1779–1801.

Green, T.C. and Figlewski, S. (1999). Market risk and model risk for a financial institution writing options, *Journal of Finance* **54**, 1465-1499.

Harrison, J. M. and Kreps, D. (1979). Martingales and arbitrage in multiperiod securities markets," *Journal of Economic Theory* **20**, 381-408.

Heston, S. (1993). A closed-form solution for options with stochastic volatility with applications to bond and currency options, *Review of Financial Studies* **6**, 327–343.

Hull, J., and White, A. (1987) The pricing of options on assets with stochastic volatilities. *Journal of Finance* **42**, 281–300.

Jackwerth, J.C. (2004). *Option-Implied Risk-Neutral Distributions and Risk Aversion,* Research Foundation of AIMR, Charlottesville, Virginia.

Jacquier, A. and Slaoui, S. (2010). Variance dispersion and correlation swaps, unpublished working paper, available at http://papers.ssrn.com/sol3/papers.cfm?abstract_id=998924

OptionMetrics, (2008). *Ivy DB File and Data Reference Manual, Version 2.5.10*.

Nelson, D. (1991). Conditional heteroskedasticity in asset returns: A new approach, *Econometrica* **59**, 347–370.

Pan, J. (2002). The jump-risk premia implicit in options: Evidence from an integrated time-series study, *Journal of Financial Economics* **63**, 3–50.

Poon, S. and Granger, C.W.J. (2003). "Forecasting volatility in financial markets: a review, *Journal of Economic Literature* **41**, 478-539.

Whaley, R. (2000). The investor fear gauge, *Journal of Portfolio Management* **26**, 12–17.

Table I: Stocks in the Data Sample

Ticker Symbol	Company Name
AA	Alcoa
AXP	American Express
BA	Boeing
BAC	Bank of America
C	Citigroup
CAT	Caterpillar
CVX	Chevron
DD	Du Pont
DIS	Disney
GE	General Electric
GM	General Motors
HD	Home Depot
HPQ	Hewlett Packard
IBM	International Business Machines
INTC	Intel
JNJ	Johnson & Johnson
JPM	JP Morgan Chase
KO	Coca Cola
MCD	McDonalds
MMM	3M CO
MRK	Merck
MSFT	Microsoft
PFE	Pfizer
T	AT&T
UTX	United Technologies
VZ	Verizon
WMT	Walmart
XOM	Exxon Mobil

Table II: EGARCH Estimates for Stock Implied Volatilities with No Common Factor

Notes: The table reports the estimated coefficients and t-statistics (in parentheses) from the 0-factor model:

$$log\left(\frac{v_{it}}{v_{it-1}}\right) = \theta_{1i} + \theta_{2i}\log(v_{it-1}) + \sigma_{it}\varepsilon_{it}$$
$$log(\sigma_{it}^2) = \phi_{1i} + \phi_{2i}|\varepsilon_{it-1}| + \phi_{3i}\varepsilon_{it-1} + \phi_{4i}log(\sigma_{it-1}^2)$$

Ticker	θ_{1i}	θ_{2i}	ϕ_{1i}	ϕ_{2i}	ϕ_{3i}	ϕ_{4i}	*t-dist d.f.*
Sample averages	-0.011 (-3.42)	-0.010 (-4.20)	-0.639 (-14.05)	0.223 (18.08)	0.070 (8.23)	0.920 (285.59)	3.88
VIX Index	0.020 (4.25)	-0.006 (-3.88)	-0.343 (-9.32)	0.086 (6.13)	0.138 (12.01)	0.952 (167.75)	7.77
AA	-0.011 (-3.87)	-0.010 (-4.19)	-0.253 (-15.53)	0.170 (18.91)	0.036 (4.74)	0.979 (430.21)	3.41
AXP	0.004 (3.63)	0.002 (2.53)	-0.153 (-14.23)	0.084 (15.13)	0.096 (19.78)	0.984 (643.38)	3.29
BA	-0.011 (-3.54)	-0.010 (-4.01)	-1.303 (-14.25)	0.306 (16.05)	0.089 (7.47)	0.827 (59.48)	4.78
BAC	0.003 (3.06)	0.001 (1.83)	-0.089 (-15.54)	0.081 (13.81)	0.082 (16.73)	0.994 (994.88)	3.82
C	-0.004 (-1.51)	-0.004 (-2.46)	-0.277 (-12.30)	0.166 (16.07)	0.066 (8.66)	0.973 (334.32)	3.66
CAT	-0.013 (-4.36)	-0.012 (-4.69)	-1.123 (-15.80)	0.329 (19.57)	0.060 (5.91)	0.859 (87.07)	3.59
CVX	-0.017 (-4.17)	-0.012 (-4.42)	-0.367 (-11.99)	0.208 (16.34)	0.025 (3.04)	0.966 (258.91)	4.06
DD	-0.005 (-2.07)	-0.004 (-2.60)	-0.112 (-7.78)	0.075 (10.01)	0.039 (7.12)	0.991 (539.03)	4.05
DIS	-0.014 (-5.20)	-0.013 (-6.27)	-1.128 (-14.19)	0.368 (19.59)	0.125 (10.46)	0.858 (73.57)	3.45
GE	-0.007 (-2.35)	-0.006 (-3.21)	-0.342 (-12.96)	0.163 (15.37)	0.076 (9.91)	0.961 (240.69)	4.44
GM	-0.008 (-3.58)	-0.008 (-4.31)	-0.334 (-16.62)	0.199 (19.09)	0.059 (9.58)	0.968 (331.63)	3.93
HD	-0.012 (-4.68)	-0.011 (-5.55)	-0.375 (-10.15)	0.128 (10.57)	0.076 (10.99)	0.953 (184.76)	4.18
HPQ	-0.010 (-4.12)	-0.014 (-6.00)	-2.016 (-23.58)	0.514 (33.60)	0.127 (10.13)	0.716 (49.82)	3.01

Table II: continued

Ticker	θ_{1i}	θ_{2i}	ϕ_{1i}	ϕ_{2i}	ϕ_{3i}	ϕ_{4i}	*t-dist d.f.*
IBM	-0.012 (-6.07)	-0.011 (-7.43)	-1.616 (-16.68)	0.310 (16.68)	0.203 (14.84)	0.760 (49.45)	3.15
INTC	-0.014 (-6.21)	-0.016 (-7.64)	-0.826 (-15.56)	0.247 (16.04)	0.129 (11.37)	0.891 (102.77)	4.23
JNJ	-0.011 (-3.09)	-0.008 (-3.64)	-0.666 (-15.06)	0.288 (20.58)	0.021 (2.03)	0.921 (138.03)	3.56
JPM	-0.006 (-2.55)	-0.006 (-3.31)	-0.232 (-21.09)	0.178 (27.67)	0.054 (8.30)	0.982 (605.28)	4.23
KO	-0.009 (-2.92)	-0.007 (-3.55)	-1.108 (-13.06)	0.335 (15.61)	0.077 (5.52)	0.858 (68.60)	4.32
MCD	-0.021 (-5.44)	-0.017 (-5.73)	-0.890 (-15.20)	0.318 (21.64)	-0.006 (-0.57)	0.890 (98.10)	4.11
MMM	-0.003 (-1.42)	-0.003 (-2.31)	-0.162 (-9.89)	0.069 (12.65)	0.096 (19.07)	0.982 (399.26)	3.61
MRK	-0.031 (-8.42)	-0.025 (-8.54)	-0.064 (-8.25)	0.059 (10.55)	-0.011 (-3.28)	0.996 (1,046.1)	3.76
MSFT	-0.008 (-3.15)	-0.009 (-4.62)	-1.654 (-17.55)	0.439 (24.37)	0.143 (12.05)	0.768 (50.03)	3.64
PFE	-0.021 (-5.58)	-0.018 (-6.02)	-0.254 (-11.69)	0.158 (17.97)	0.000 (0.07)	0.976 (299.09)	4.04
T	-0.011 (-3.72)	-0.010 (-4.67)	-0.286 (-13.11)	0.184 (30.82)	0.039 (6.33)	0.973 (283.91)	3.75
UTX	-0.012 (-4.24)	-0.010 (-4.98)	-0.768 (-13.81)	0.213 (14.66)	0.104 (10.89)	0.902 (111.49)	3.25
VZ	-0.008 (-2.40)	-0.007 (-3.06)	-0.793 (-19.79)	0.341 (27.81)	0.061 (7.02)	0.908 (159.04)	3.17
WMT	-0.005 (-1.80)	-0.005 (-2.37)	-0.323 (-8.60)	0.133 (12.18)	0.066 (8.26)	0.963 (177.54)	4.75
XOM	-0.024 (-6.08)	-0.017 (-6.43)	-0.366 (-9.17)	0.182 (12.99)	0.040 (4.01)	0.963 (180.16)	5.49

Table III: EGARCH Estimates for Stock Implied Volatilities with One Common Factor

Notes: The table reports the estimated coefficients and t-statistics (in parentheses) from the 1-factor model:

$$log\left(\frac{v_{it}}{v_{it-1}}\right) = \theta_{1i} + \theta_{2i}\log\left(\frac{vix_t}{vix_{t-1}}\right) + \theta_{3i}\log(v_{it-1}) + \theta_{4i}\log(vix_{t-1}) + \sigma_{it}\varepsilon_{it}$$

$$log(\sigma_{it}^2) = \phi_{1i} + \phi_{2i}|\varepsilon_{it-1}| + \phi_{3i}\varepsilon_{it-1} + \phi_{4i}log(\sigma_{it-1}^2)$$

Ticker	θ_{1i}	θ_{2i}	θ_{3i}	θ_{4i}	ϕ_{1i}	ϕ_{2i}	ϕ_{3i}	ϕ_{4i}	*t-dist d.f.*
Sample averages	-0.082 (-5.86)	0.325 (31.58)	-0.022 (-6.25)	0.018 (5.37)	-1.351 (-6.14)	0.299 (8.24)	0.043 (1.90)	0.816 (74.56)	3.98
AA	-0.066 (-6.34)	0.205 (21.34)	-0.021 (-6.63)	0.014 (5.53)	-0.112 (-4.32)	0.090 (5.72)	0.009 (.75)	0.992 (311.96)	3.35
AXP	-0.084 (-5.61)	0.359 (37.87)	-0.018 (-5.59)	0.021 (5.51)	-0.433 (-6.15)	0.217 (8.20)	0.053 (3.14)	0.954 (97.46)	3.41
BA	-0.130 (-8.13)	0.308 (31.84)	-0.040 (-8.58)	0.028 (7.64)	-3.020 (-6.11)	0.379 (8.21)	0.024 (.80)	0.572 (7.61)	4.64
BAC	-0.084 (-5.39)	0.362 (30.59)	-0.015 (-5.12)	0.022 (5.30)	-0.277 (-6.52)	0.172 (8.33)	0.042 (2.67)	0.974 (163.66)	3.84
C	-0.102 (-5.34)	0.422 (37.15)	-0.021 (-5.56)	0.026 (5.15)	-0.203 (-5.46)	0.140 (7.55)	0.047 (3.33)	0.983 (199.53)	4.14
CAT	-0.070 (-5.20)	0.250 (28.16)	-0.025 (-5.92)	0.014 (4.54)	-1.624 (-5.80)	0.340 (8.27)	0.020 (.79)	0.786 (19.22)	3.76
CVX	-0.053 (-4.88)	0.238 (26.14)	-0.019 (-5.64)	0.009 (3.66)	-0.316 (-5.41)	0.174 (7.95)	-0.001 (-.06)	0.971 (127.20)	4.69
DD	-0.116 (-7.35)	0.320 (33.80)	-0.031 (-7.61)	0.025 (7.04)	-3.456 (-5.84)	0.364 (7.37)	0.028 (.85)	0.503 (5.57)	4.23
DIS	-0.097 (-6.90)	0.278 (27.65)	-0.025 (-7.02)	0.022 (6.61)	-1.853 (-7.24)	0.427 (9.44)	0.094 (3.41)	0.750 (19.36)	3.95
GE	-0.132 (-6.89)	0.526 (47.90)	-0.029 (-7.12)	0.031 (6.62)	-1.148 (-7.43)	0.364 (10.10)	0.014 (.57)	0.856 (36.88)	4.48
GM	-0.028 (-3.57)	0.270 (26.43)	-0.008 (-4.02)	0.006 (2.84)	-0.287 (-6.16)	0.173 (7.74)	0.036 (2.27)	0.973 (153.59)	3.68
HD	-0.116 (-7.42)	0.378 (37.06)	-0.031 (-7.54)	0.027 (7.21)	-2.417 (-5.42)	0.340 (7.87)	0.060 (2.19)	0.655 (9.53)	4.29
HPQ	-0.028 (-3.15)	0.290 (28.78)	-0.010 (-3.87)	0.007 (2.88)	-2.083 (-7.13)	0.455 (9.52)	0.082 (2.78)	0.698 (15.02)	3.02

Table III: continued

Ticker	θ_{1i}	θ_{2i}	θ_{3i}	θ_{4i}	ϕ_{1i}	ϕ_{2i}	ϕ_{3i}	ϕ_{4i}	*t-dist d.f.*
IBM	-0.033	0.388	-0.010	0.007	-2.626	0.360	0.136	0.605	3.18
	(-2.77)	(38.60)	(-3.62)	(2.42)	(-5.75)	(7.34)	(4.38)	(8.23)	
INTC	-0.076	0.391	-0.026	0.017	-1.421	0.318	0.081	0.806	4.00
	(-6.43)	(36.76)	(-7.49)	(5.85)	(-5.62)	(8.12)	(3.32)	(20.91)	
JNJ	-0.070	0.290	-0.019	0.014	-1.517	0.442	0.029	0.793	3.54
	(-5.77)	(25.75)	(-6.33)	(4.90)	(-7.63)	(9.72)	(1.02)	(25.20)	
JPM	-0.130	0.397	-0.027	0.033	-0.214	0.149	0.063	0.983	4.53
	(-7.14)	(35.93)	(-7.12)	(7.03)	(-6.41)	(8.26)	(4.80)	(226.24)	
KO	-0.072	0.340	-0.017	0.016	-2.190	0.391	0.055	0.694	4.24
	(-5.31)	(33.00)	(-5.71)	(4.87)	(-6.69)	(8.79)	(1.87)	(13.72)	
MCD	-0.097	0.257	-0.030	0.019	-1.287	0.362	0.000	0.832	4.17
	(-7.71)	(23.60)	(-7.85)	(6.88)	(-6.51)	(9.24)	(.00)	(27.84)	
MMM	-0.087	0.313	-0.026	0.017	-1.969	0.386	0.038	0.729	3.64
	(-5.88)	(33.48)	(-6.40)	(5.36)	(-6.49)	(8.37)	(1.32)	(15.95)	
MRK	-0.089	0.306	-0.026	0.018	-0.098	0.063	0.024	0.991	3.75
	(-6.49)	(29.99)	(-6.61)	(5.86)	(-4.01)	(5.22)	(2.55)	(305.57)	
MSFT	-0.068	0.399	-0.019	0.016	-2.595	0.465	0.128	0.621	3.80
	(-5.11)	(37.51)	(-6.11)	(4.75)	(-7.49)	(10.32)	(4.36)	(11.16)	
PFE	-0.073	0.337	-0.023	0.014	-0.327	0.147	0.007	0.962	3.84
	(-5.63)	(29.31)	(-6.22)	(4.77)	(-4.61)	(6.39)	(.44)	(93.35)	
T	-0.080	0.248	-0.019	0.019	-1.246	0.419	-0.006	0.839	3.96
	(-5.82)	(22.26)	(-6.00)	(5.46)	(-7.70)	(10.92)	(-.22)	(33.27)	
UTX	-0.092	0.261	-0.025	0.019	-1.670	0.375	0.064	0.777	3.42
	(-6.69)	(29.14)	(-7.22)	(6.15)	(-6.46)	(8.19)	(2.41)	(20.26)	
VZ	-0.075	0.232	-0.019	0.017	-1.544	0.387	0.036	0.790	3.43
	(-5.57)	(23.27)	(-5.69)	(5.26)	(-7.28)	(9.32)	(1.37)	(24.03)	
WMT	-0.070	0.370	-0.018	0.016	-1.469	0.307	0.030	0.801	4.77
	(-5.32)	(34.64)	(-5.41)	(5.03)	(-5.51)	(7.58)	(1.20)	(20.01)	
XOM	-0.083	0.369	-0.026	0.015	-0.432	0.161	0.014	0.952	5.80
	(-6.22)	(36.39)	(-6.91)	(5.12)	(-4.69)	(6.72)	(.89)	(75.38)	

Table IV: Hedging Vega Risk for an Individual Stock with the Most Highly Correlated 1, 2, or 3 Stocks

The Table presents results of vega hedging each individual stock's implied volatility changes using linear payoff derivative contracts tied to the VIX index and to the IVs of one other single stock. Hedge performance is the standard deviation of the change in the value of a portfolio consisting of a linear volatility contract on a single name by itself, on the stock's IV hedged with only VIX contracts, and on an "extended" hedge using VIX contracts plus volatility contracts on the other stock in the sample with the most highly correlated idiosyncratic IV risk exposures on each date. Columns 3 and 4 show, respectively, the average t-statistic across names on the difference in mean squared error between the unhedged position and a hedged position, and between the extended hedge and the hedge with only the VIX. Columns 5 and 6 report t-statistics on the difference in mean squared error for the pooled sample, using robust standard errors that are corrected for heteroskedasticity across names and contemporaneous correlation in the errors.

		t-statistics on the Difference in Mean Squared Hedging Error			
		Average across Individual Names		Full sample with Robust Standard Errors	
# hedging stocks	Hedge performance	Hedge vs Unhedged	Extended hedge vs VIX only	Hedge vs Unhedged	Extended hedge vs VIX only
No hedge	0.0227	--	--	--	--
VIX only	0.01971	7.301	--	8.165	--
VIX +1	0.01970	3.304	0.982	3.101	0.173

Table V: Hedging Vega Risk for an Individual Stock with Stocks from the Same Industry

The Table presents results of vega hedging an individual stock's implied volatility changes using linear payoff derivative contracts tied the VIX index and to the IVs of one or two other stocks from the same 4-digit SIC industry classification. Hedge performance in columns 5-7 is, respectively, the standard deviation of the change in the value of a portfolio consisting of a linear volatility contract on a single stock by itself, on the stock hedged with only VIX contracts, and on an "extended" hedge using VIX contracts and volatility contracts on the hedging stocks. Columns 8 and 9 report, respectively, the t-statistics on the difference in mean squared error between the unhedged position and a hedged position, and between the extended hedge and the hedge with only the VIX.

Target stock	SIC Code	Hedging stocks		Hedge performance			t-statistics on the difference in squared errors	
				Unhedged	VIX only hedge	Extended hedge	Extended hedge vs Unhedged	Extended hedge vs VIX only
MRK	2834	JNJ	PFE	0.0162	0.0142	0.0133	6.502	3.278
PFE	2834	JNJ	MRK	0.0184	0.0162	0.0156	6.238	1.879
JNJ	2834	MRK	PFE	0.0160	0.0145	0.0135	4.727	2.623
BAC	6021	C	JPM	0.0515	0.0474	0.0393	4.247	3.405
C	6021	BAC	JPM	0.0421	0.0367	0.0342	1.938	0.727
JPM	6021	BAC	C	0.0347	0.0290	0.0221	5.212	3.842
CVX	2911	XOM	--	0.0165	0.0141	0.0097	5.451	4.892
XOM	2911	CVX	--	0.0171	0.0136	0.0103	5.998	5.111
HPQ	3570	IBM	--	0.0239	0.0218	0.0224	1.883	-0.676
IBM	3570	HPQ	--	0.0199	0.0168	0.0182	1.504	-1.159
T	4813	VZ	--	0.0187	0.0169	0.0170	1.012	-0.031
VZ	4813	T	--	0.0193	0.0178	0.0225	-0.810	-1.154

Figure 1: Implied Volatilities of 28 Stocks and the VIX Index, 1996 - 2009

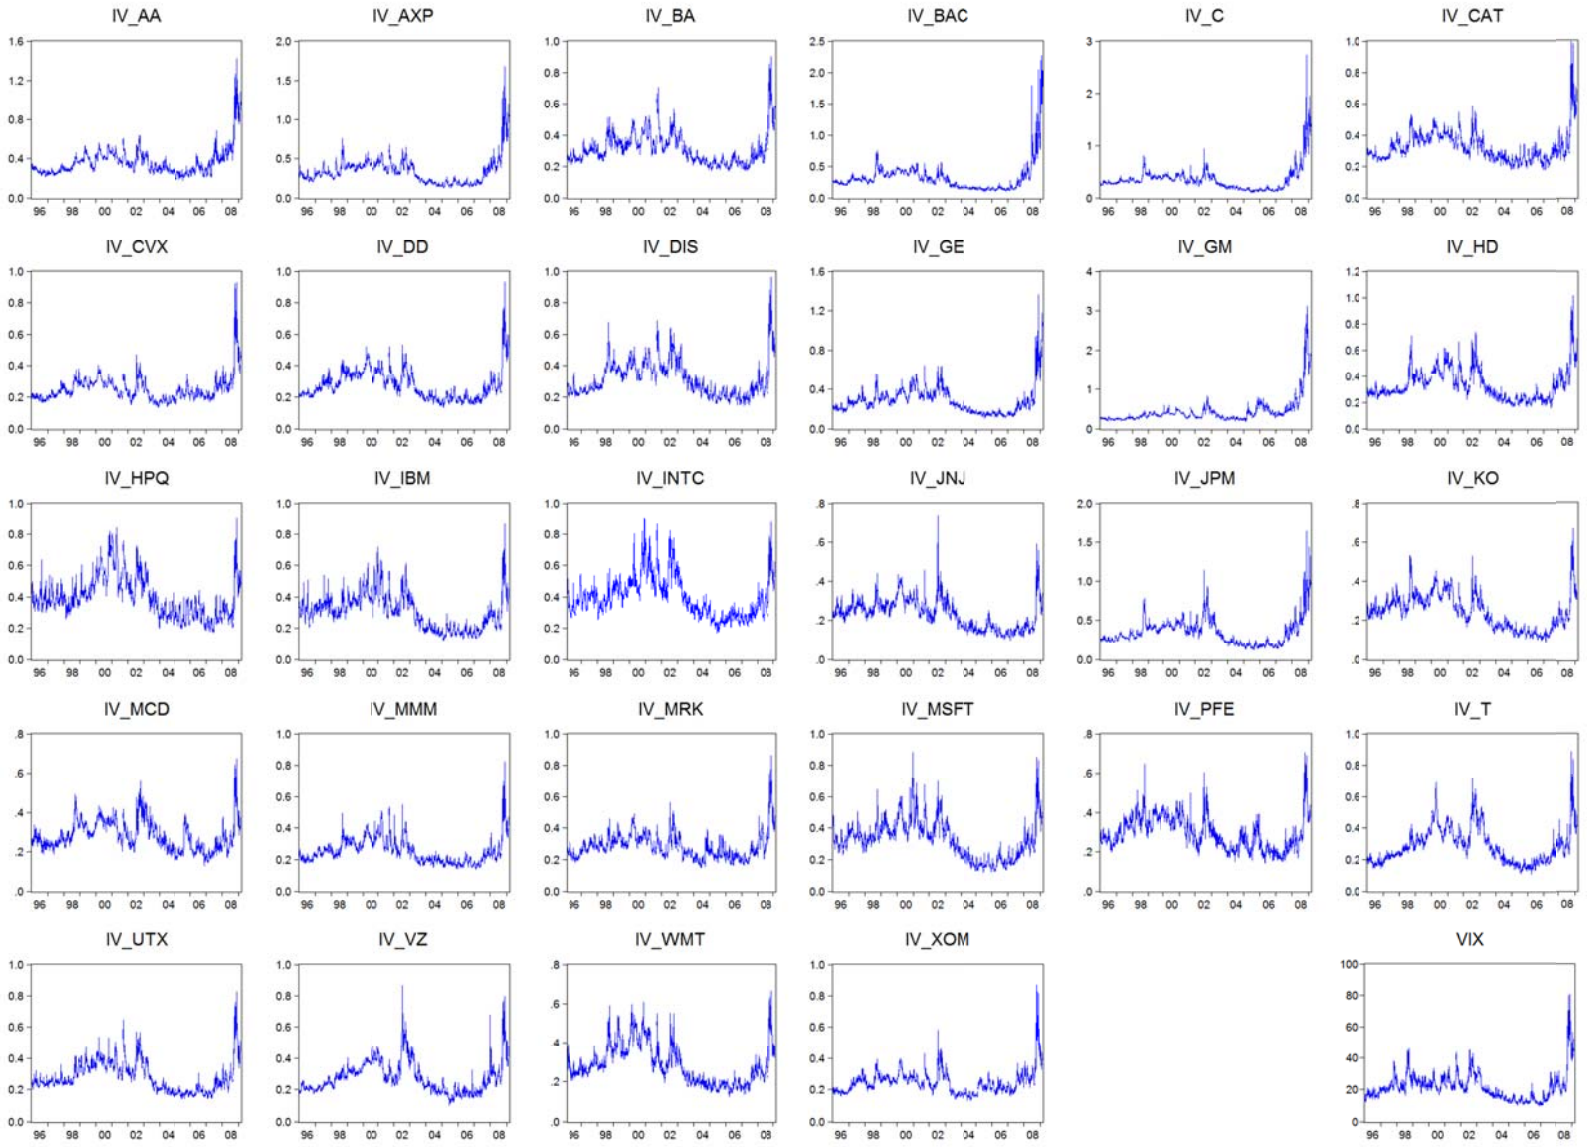


Figure 2: Mean Estimated Error Correlation Across Pairs

