Probabilistic Interpretation of Black Implied Volatility

Peter Carr, Liuren Wu, Yuzhao Zhang
pc73@nyu.edu

Initial version: September 15, 2018
Current version: August 31, 2021
File reference: Probabilistic Interpretation of BSIV1.tex

Abstract

We use a market model of implied volatility to develop an implied volatility smile. The implied variance rate is given a simple probabilistic representation.

1 Assumptions

For some fixed option maturity date T, assume zero interest rates and zero dividends for a stock over the time period $[0, T]$. Further assume that the spot price process S for the underlying stock index is a strictly positive local martingale under a risk-neutral probability measure $\mathbb{Q}$. Under these conditions, the results of Merton can be used to produce arbitrage free ranges for option prices and hence implied volatilities. The restrictions on option prices are fairly weak and simple to state geometrically. For example, call prices must lie between intrinsic value and stock price and must be decreasing convex functions of strike price. The corresponding restrictions on implied volatility are also fairly weak, but difficult to state geometrically. For example, implied volatility can be concave in strike price but not too concave.

In an effort to reduce the range of arbitrage-free implied volatilities and to simplify their geometric properties, dynamical restrictions on assets prices can be imposed. For example, one can further require that the stock price never jumps. In this case, there exists a bounded stochastic process σ_t, called the instantaneous volatility of the stock, such that S_t is the unique solution of the following stochastic differential equation (SDE):

$$dS_t = \sigma_t S_t dW_t, \qquad t \geq 0,$$

where W is standard Brownian motion (SBM) under $\mathbb{Q}$. We allow the increments of σ to be correlated with those of S, but we do not directly specify the σ process. Instead, we will partially specify the risk-neutral dynamics of Black implied volatility (IV) across a continuum of positive strike prices at one fixed maturity date. When coupled with the stock price dynamics, the assumed dynamics for implied volatilities restrict the set of arbitrage-free dynamics for the option prices.

For our fixed maturity date T, let $I_t(K)$ be the Black IV at strike price $K > 0$. Suppose that three or more IV's are known. The goal is to develop an IV curve $I_t(K), K > 0$, which has five properties:

1. When exactly 3 IV's are given, $I_t(K), K > 0$ interpolates the three given implied volatilities.
2. For $K > 0$, $k \equiv \ln(K/S)$, and fixed S, $v_t^2(k) \equiv I_t^2(K)$ is a simple explicit algebraic function of k.
3. For $k \in \mathbb{R}$, $v_t^2(k)$ has a simple probabilistic interpretation.
4. $I_t(K), K > 0$ is free of factor-based arbitrage.
5. For $k \in \mathbb{R}$, $v_t^2(k)$ is testable, i.e. it can be rejected using data.

The third desired property is the main contribution of this paper. The graph of implied variance rates across moneyness is well-defined, widely studied, but not well understood. In fact, the authors of this paper have never seen a probabilistic interpretation of all implied variance rates at a given maturity. At-the money (ATM) IV has been argued to be the best forecast of subsequent realized volatility. A convex combination of all of the implied variance rates has been argued to be the best forecast of the subsequent realized variance rate. The slope and curvature of ATM IV in moneyness have been argued to say something about skewness and excess kurtosis of the underlying stock index. However, option traders and options quants would both be hard pressed to provide an exact direct probabilistic interpretation of each of the uncountably infinite number of IV's at each strike price and maturity date that an arbitrage-free model produces.

The fourth desired property of an IV smile is the absence of factor-based arbitrage. Factor-based arbitrage could potentially arises as a consequence of an assumption that all co-terminal IV's are continuous semi-martingales driven by a single factor called Z. Z is assumed to be an SBM under risk-neutral measure $\mathbb{Q}$. We let the increments of Z be correlated with those of S. The correlation coefficient is an unspecified stochastic process $\rho_t \in [-1, 1]$ at each $t \in [0, T]$. Thus all option prices are driven by just two SBM's W and Z. At each time $t \in [0, T]$, the IV's need to be set such that the instantaneous profit from setting up a delta-neutral and vega-neutral portfolio is always zero. Put another way, all option prices must be $\mathbb{Q}$ local martingales over $[0, T]$. For $t \in [0, T]$, let $C_t(K)$ be the market price of a call and for $S > 0, \sigma > 0$, let $B(S, \sigma, t)$ be the Black call pricing formula, where the strike price K and maturity date T are suppressed for notational simplicity:

$$B(S, \sigma, t) = SN\left(\frac{\ell}{\sigma\sqrt{T-t}} + \frac{\sigma\sqrt{T-t}}{2}\right) - KN\left(\frac{\ell}{\sigma\sqrt{T-t}} - \frac{\sigma\sqrt{T-t}}{2}\right), \quad (1)$$

where $\ell \equiv -k = \ln(S/K)$ and as usual, $N(z) \equiv \int_{-\infty}^{z} \frac{e^{-\frac{x^2}{2}}}{\sqrt{2\pi}} dx$ is the standard normal cumulative distribution function. By the definition of IV, $C_t(K) = B(S_t, I_t(K), t)$ for all $t \in [0, T]$ and for all $K > 0$. Under $\mathbb{Q}$ and for fixed $T > 0$, we assume that $I_t(K)$ is the unique solution of the following SDE:

$$dI_t(K) = \mu_t(I_t(K))I_t(K)dt + \omega_t I_t(K)dZ_t, \qquad t \geq 0, K > 0. \quad (2)$$

The vol vol ω is assumed to be the same stochastic process across all strike prices $K > 0$. Like the instantaneous volatility σ_t and the correlation coefficient ρ_t, the vol vol ω_t does not depend on strike price $K > 0$ or the IV $I(K) > 0$. In contrast, the proportional risk-neutral drift coefficient μ_t is allowed to depend on $I_t(K)$, as indicated in (2). We assume that this dependence is such that IV's are always strictly positive, i.e. $I_t(K) > 0$ for all $t \in [0, T]$. For $t \in [0, T]$, let $\gamma_t \equiv \sigma_t \rho_t \omega_t$ be the stochastic process capturing the covariation rate between $\ln S$ and $\ln I$. Since σ_t, ρ_t, and ω_t do not depend on K or $I_t(K)$, neither does the covariation rate γ_t. The four stochastic processes $\sigma_t, \gamma_t, \omega_t$, and $\mu_t(I_t(K))$ are all unspecified, but must evolve in a way consistent with the dynamics in (2) governing IV. For example, the four stochastic processes should have continuous sample paths because the results of our analysis will imply that a jump in any of the 4 processes would induce a jump in each IV. Recall that we are assuming that we only observe three IV's. To obtain a unique IV smile, we will choose the proportional risk-neutral drift coefficient μ_t to depend on $\sigma_t, \gamma_t, \omega_t$, and $I_t(K)$ such that our arbitrage-free IV smile has a simple probabilistic interpretation. The particular restriction that we impose is:

$$\mu_t(I_t(K)) = \frac{\omega_t^2}{8} I_t^2(K) - \frac{\gamma_t}{2}, \qquad t \in [0, T], K > 0. \quad (3)$$

The fifth desired property is advantageous on the view that the input data is clean and that the option market is sometimes inefficient when quoting option prices or IV's. When the factor-arbitrage-free family of IV curves does not fit the market IV quotes, one can either take the view that our assumed factor model is wrong, or else take the view that market option prices are failing to reflect the actual known factor structure. On the latter view, portfolios can be formed which at least have positive risk-neutral drift and which at best also zero out the martingale components of P&L.

2 Probabilistic Interpretation of the Implied Variance Rate

In this section, we develop a probabilistic interpretation of the implied variance rate, I_t^2 at each positive strike price $K > 0$ and for fixed maturity date $T > 0$. In a world of zero interest rates and zero dividends before T, the implied variance rate $I_t^2(K), K > 0, t \in [0,T]$ can be defined as the ratio of negated theta to cash gamma. Cash gamma is defined as the product of the squared stock price, i.e. S^2 and the second partial derivative of the option price with respect to S, i.e. gamma. As its name suggests, cash gamma has the same units as the option premium, eg. dollars. Cash gamma in the Black model depends on $S > 0, \sigma > 0$, and $t \in [0,T]$ as follows:

$$S^2 B_{11}(S,\sigma,t) = \frac{K}{\sigma\sqrt{T-t}} N'\left(\frac{\ell}{\sigma\sqrt{T-t}} - \frac{\sigma\sqrt{T-t}}{2}\right), \tag{4}$$

where subscripts of B denote partial derivatives and recall $\ell \equiv \ln(S/K)$. Still in the Black model, define cash vega as $\sigma B_2(S,\sigma,t)$, cash vanna as $S\sigma B_{12}(S,\sigma,t)$, and cash volga as $\sigma^2 B_{22}(S,\sigma,t)$. All 3 of these cash greeks have the same units as cash gamma, eg. dollars, and moreover are simply the product of a dimensionless quantity and cash gamma:

$$\sigma B_2(S,\sigma,t) = \sigma^2(T-t) \times S^2 B_{11}(S,\sigma,t) \tag{5}$$

$$\sigma S B_{12}(S,\sigma,t) = (k + \sigma^2(T-t)/2) \times S^2 B_{11}(S,\sigma,t) \tag{6}$$

$$\sigma^2 B_{22}(S,\sigma,t) = (k^2 - \sigma^4(T-t)^2/4) \times S^2 B_{11}(S,\sigma,t). \tag{7}$$

Zeroing out the risk-neutral drift of a call's price implies that the call's theta must negate the P&L contributions from cash vega, cash gamma, cash vanna, and cash volga. As a result, the absence of factor arbitrage implies that negated theta is an inner product of a four vector of dynamical coefficients governing proportional changes, $[\mu_t(I_t(K)), \frac{\sigma_t^2}{2}, \gamma_t, \frac{\omega_t^2}{2}]^T$ with a four vector of cash greeks $[I_t(K)B_2, S_t^2 B_{11}, S_t I_t(K) B_{12}, I_t^2(K) B_{22}]^T$:

$$-B_3 = \mu_t(I_t(K)) I_t(K) B_2 + \frac{\sigma_t^2}{2} S_t^2 B_{11} + \gamma_t S_t I_t(K) B_{12} + \frac{\omega_t^2}{2} I_t^2(K) B_{22}. \tag{8}$$

Dividing (8) by halved cash gamma, $S^2 B_{11}/2 > 0$, and using (5) to (7) implies:

$$I_t^2(K) = 2\mu_t(I_t(K)) I_t^2(K)(T-t) + \sigma_t^2 + 2\gamma_t(k + I_t^2(K)(T-t)/2) + \omega_t^2(k^2 - I_t^4(K)(T-t)^2/4). \tag{9}$$

Substituting (3) in (9) simplifies it so that the implied variance rate $v_t^2(k) \equiv I_t^2(K)$ is quadratic in $k \equiv \ln(K/S)$:

$$v_t^2(k) = \sigma_t^2 + 2\gamma_t k + \omega_t^2 k^2, \qquad k \in \mathbb{R}, \tag{10}$$

where recall $\gamma_t \equiv \sigma_t \rho_t \omega_t$. This is the simple explicit formula for the implied variance rate that we seek. Since $|\rho_t| \leq 1$, this quadratic function of k is positive everywhere. Since $\omega_t^2 > 0$ our quadratic function must open upwards. It follows that under our assumptions, IV's must smile everywhere and can never frown anywhere. When no dynamical restrictions are imposed, IV's can frown and yet be free of arbitrage. For example, if instead of our continuous time dynamics, the final stock price S_T takes only 2 values, $S_0 u$ for $u > 1$ and S_0/u, then the absence of arbitrage leads to negative curvature of IV in strike price K at every $K \in (S_0/u, S_0 u)$ (and for every $k \in (\ln S_0 - \ln u, \ln S_0 + \ln u)$. In contrast, under our market model, the absence of arbitrage leads positive curvature at every strike price and time.

Our arbitrage-free implied variance rate smile (10) can also be expressed as:

$$v_t^2(k)dt = (\sigma_t dW_t + k\omega_t dZ_t)^2 . \tag{11}$$

This is the simple probabilistic interpretation of the implied variance rate that we seek. When $k = 0$, then the at-the-money implied variance rate is the instantaneous variance rate of the underlying stock return, since $\frac{dS_t}{S_t} = \sigma_t dW_t$. When $k \neq 0$, then the away from-the-money Black implied variance rate is a variance rate, but not just of the underlying stock return. Rather the away from-the-money Black implied variance rate is the variance rate of a linear combination of the increments of the two SBM's W and Z driving option prices. The further away from the money the option is, the greater the weight $k\omega_t$ on the SBM Z driving IV's.

Our simple explicit formula (10) for the implied variance rate can be tested empirically, so long as three or more IV's are given by the market. When exactly 3 IV's are given, one can always fit a quadratic function to them. Notice that the implied variance rate in (10) is linear in σ_t^2, γ_t and ω_t^2. Given three option IV's, one can square them to get implied variance rates and then write (10) thrice with these squared values on its LHS. One can then invert the 3x3 linear system to get the coefficients σ_t^2, γ_t, and ω_t^2. From here, it is straightforward to get $\sigma_t = \sqrt{\sigma_t^2}$, $\omega_t = \sqrt{\omega_t^2}$, and $\rho_t = \frac{\gamma_t}{\sigma_t \omega_t}$.

When our assumptions are not necessarily holding, the quadratic function of k that one fits to 3 given IV quotes need not lie in the restricted class of quadratic functions described by (10). Consider an arbitrary quadratic function of moneyness k:

$$q(k) = c_0 + c_1 k + c_2 k^2, k \in \mathbb{R},$$

where the constant coefficients c_0, c_1, and c_2 are real-valued. Our restricted class of quadratics imposes a single restriction on each of the three coefficients:

1. $c_0 > 0$
2. $c_2 > 0$
3. $\frac{c_1}{2\sqrt{c_0}\sqrt{c_2}} \in [-1, 1]$

The last condition can be interpreted as requiring that the implied correlation $\rho_t \in [-1, 1]$. Straightforward calculus show thats when the three restrictions above apply, the minimum of our quadratic function occurs at $k = -\frac{\sigma_t \rho_t}{\omega_t}$, i.e. the negative of the slope coefficient in a regression of $\ln S$ on $\ln I$. At this minimum, the implied variance rate $v_t^2(k)$ is positive if the three conditions are holding.

A rejection of any of these three restrictions can be regarded as either a failure of the (S, I) factor model or a failure of the options market maker. In the latter case, any of the three rejections leads to a factor-based arbitrage. Rejection of the first condition leads to vol trading. Rejection of the second condition leads to smile trading. Rejection of the third condition leads to skew trading.

To test this model when four or more IV's are given, one can do a linear regression at a fixed spot $S > 0$ and time $t \in [0, T]$, of the implied variance rates $v_t^2(k)$ on an intercept, on varying moneyness $k \equiv \ln(K/S)$, and on varying squared moneyness k^2. If the model works perfectly, then the R^2 from this linear regression should be 100%. Moreover, the above three conditions on the coefficients c_0, c_1, and c_2 must be always respected.

When three or more IV quotes are given by the market, a further test of the factor model is to see if the three coefficients produced by the quadratic fit or linear regression capture the realized variance of returns on the underlying, the realized covariation of percentages changes in stock price and any IV, and finally the realized variance of percentage changes in (all co-terminal) IV's.

3 Financial Interpretations of the Implied Variance Rate

In this section, we provide two financial interpretations of the implied variance rate. The first uses a dynamic trading strategy in a call option with a fixed strike price. The second uses a dynamic trading strategy in an always at-the-money (ATM) call option.

3.1 Fixed Strike Price

Consider the gains that arise from a dynamic trading strategy that always holds N_t^c calls, delta-hedged with futures. This strategy is not self-financing, but a trivial extension of the portfolio can be forced to be self-financing by introducing a riskless asset, which has zero return by assumption. The position in the riskless asset is used to both self-finance and carry the gains from the calls and futures through time, that are not carried by the call position itself. Let $gN_t^c C_t(K)$ denote the gain at time $t \in [0, T]$ from always holding the N_t^c delta-hedged calls. From Itô's formula and financial considerations, we have $gN_t^c C_t(K) = N_t^c B_2(S_t, I_t(K), t) I_t(K) \omega_t dZ_t$, since the arbitrage-free gains from any dynamic trading strategy must be a $\mathbb{Q}$ local martingale under zero interest rates, and the contributions from W have been delta-hedged away Now suppose we set N_t^c to the ratio of moneyness to cash vega, i.e.:

$$N_t^c \equiv \frac{k}{I_t(K) B_2(S_t, I_t(K), t)}.$$

Then we have $gN_t^c C_t(K) = \frac{k}{I_t(K) B_2(S_t, I_t(K), t)} B_2(S_t, I_t(K), t) I_t(K) \omega_t dZ_t = k \omega_t dZ_t$.

Suppose that an investor supplements the above dynamic trading strategy by furthermore raising the exposure to futures by one dollar. This is equivalent to always keeping one dollar invested in the stock and financing this dynamic share holding via borrowing and lending. Let V_t be the value of the combined dynamic trading strategies at time $t \in [0, T]$. Then the gain from the overall dynamic trading strategy is:

$$gV_t = \frac{dS_t}{S_t} + gN_t^c C_t(K) = \sigma_t dW_t + k\omega_t dZ_t.$$

Now recall from (11) that:

$$v_t^2(k)dt = (\sigma_t dW_t + k\omega_t dZ_t)^2. \tag{12}$$

Thus, the implied variance rate $v_t^2(k)$ captures the rate at which the quadratic variation of the portfolio value grows through time:

$$v_t^2(k)dt = \left(\frac{dS_t}{S_t} + gN_t^c C_t(K)\right)^2, \tag{13}$$

where $N_t^c \equiv \frac{k}{I_t(K)B_2(S_t, I_t(K), t)}$, $t \in [0, T]$. It is interesting to note that S_t is also the value at time $t \in [0, T]$ of a fixed strike call since $S_t = C_t(0)$ under our zero dividends assumption.

3.2 Floating Strike Price

Suppose that the strike price of a holding of one call is allowed to move through time. Letting K_t denote the strike price of the call held at time $t \in [0, T]$, the value of the call holding at time $t \in [0, T]$ is:

$$B(S_t, I_t(K_t), t; K_t, T) = S_t N\left(\frac{\ln(K_t/S_t)}{I_t(K_t)\sqrt{T-t}} + \frac{I_t(K_t)\sqrt{T-t}}{2}\right) - K_t N\left(\frac{\ln(K_t/S_t)}{I_t(K_t)\sqrt{T-t}} - \frac{I_t(K_t)\sqrt{T-t}}{2}\right). \tag{14}$$

Although only one call is being held at each time $t \in [0, T]$, a dynamic trading strategy is needed to effect this holding as the strike moves through time. The cash vega of the unit holding of a call with a floating strike K_t is given by:

$$I_t(K_t)B_2(S_t, I_t(K_t), t; K_t, T) = I_t(K_t)K_t\sqrt{T-t}N'\left(\frac{\ln(K_t/S_t)}{I_t(K_t)\sqrt{T-t}} - \frac{I_t(K_t)\sqrt{T-t}}{2}\right). \tag{15}$$

Now consider the cash value and cash vega when $K_t = S_t$, so that the one call being held at each time $t \in [0, T]$ is always at-the-money (ATM). Let $A_t \equiv A(S_t, I_t(S_t), t; S_t, T)$ be the simpler special case of call value when $K_t = S_t$:

$$A_t = A(S_t, I_t(S_t), t; S_t, T) = S_t\left[N\left(\frac{I_t(S_t)\sqrt{T-t}}{2}\right) - N\left(-\frac{I_t(S_t)\sqrt{T-t}}{2}\right)\right]. \tag{16}$$

The cash vega of an always ATM ($K_t = S_t$) call also simplifies to:

$$I_t(S_t)B_2(S_t, I_t(S_t), t; S_t, T) = I_t(S_t)S_t\sqrt{T-t}N'\left(-\frac{I_t(S_t)\sqrt{T-t}}{2}\right). \tag{17}$$

Consider the gains from always holding one ATM call which is also delta hedged:

$$gA_t = A_2(S_t, I_t(S_t), t; S_t, T)I_t(S_t)\omega_t dZ_t, \qquad t \in [0, T]. \tag{18}$$

If we change the scale of the holdings from one to the reciprocal of cash vega, then:

$$\frac{gA_t}{A_2(S_t, I_t(S_t), t; S_t, T)I_t(S_t)} = \omega_t dZ_t, \qquad t \in [0, T]. \tag{19}$$

Suppose we wish to financially interpret the implied variance rate $v_t^2(k)$ of an option with moneyness $k = \ln(K/S)$. Consider the gains that arise from keeping one dollar in stock and holding the ratio of moneyness k to ATM cash vega in always ATM calls. Thus, the implied variance rate $v_t^2(k)$ captures the rate at which the quadratic variation of this portfolio value grows through time:

$$v_t^2(k)dt = \left(\frac{dS_t}{S_t} + k\frac{gA_t}{A_2(S_t, I_t(S_t), t; S_t, T)I_t(S_t)}\right)^2, \qquad t \in [0, T]. \tag{20}$$

4 Approximation for Short Maturity Options

It is well known that in the Black model, the value of an ATM call is roughly proportional to its IV, especially at short maturities. As a result, the cash vega of an ATM call is well approximated by its cash value. Hence, the quantity $\frac{gA_t}{A_2(S_t,I_t(S_t),t;S_t,T)I_t(S_t)}$ in (20) can be approximated by the return $\frac{gA_t}{A_t}$ from always keeping one dollar in an always ATM call. In fact, when the always ATM call is always delta-hedged, the hedge portfolio return can be shown to be:

$$\frac{gA_t}{A_t} = e^{-I_t^2(S_t)(T-t)/8}\omega_t dZ_t, \qquad t \in [0, T]. \tag{21}$$

Suppose that the total ATM implied variance remaining, $I_t^2(S_t)(T-t)$, is small, either because the ATM implied variance rate $I_t^2(S_t)$ is small, or the time to maturity, $T-t$, is small, or both. Then from (21), we have $\frac{gA_t}{A_t} \approx \omega_t dZ_t$, and hence from (20):

$$v_t^2(k)dt \approx \left(\frac{dS_t}{S_t} + k\frac{gA_t}{A_t}\right)^2, \qquad t \in [0, T]. \tag{22}$$

In words, the implied variance rate at time $t \in [0, T]$ of an option with moneyness $k \equiv \ln(K/S_t)$ is approximately the rate at which the quadratic variation of a portfolio's value is growing at t, where the portfolio has k times as many dollars invested in always ATM calls as it has in the underlying stock. The approximation error approaches zero as the time to maturity decreases.

5 Empirical Tests of the Implied Variance Rate Formula

In this section, we use S&P500 index options of various maturities to provide an empirical test of our results. The sample is from 1996 to 2017, comprising 5,537 trading days. For each of the 5,537 trading days and for each maturity, there were more than 3 strike prices so a cross sectional linear regression was run. Over the 5537 trading days, we have a total of 73097 cross-sectional linear regressions, about 13 maturities per day.

Using this data we express our findings using six figures and three tables. The next three figures show the time series of the 3 slope coefficients in a linear regression on intercept, $2k$, and k^2. The intercept is always positive, as predicted by the model. Moreover, spikes in the intercept occur at known times of high volatility, suggesting the intercept is an alternative to VIX^2 as a forecast of the subsequent realized variance rate of returns on S&P500. The slope coefficient on k^2 is positive on 94% of the observations, as predicted by the model. The 6% of the observations for which the slope coefficient on k^2 is negative are days and maturities for which the least squares fit is a frown. At such times, either the factor model is unsupported or else a smile trade is available. Further empirical work on the profitability of such a smile trade is required. The third figure shows that the implied covariation rate is always negative. This means that the ATM slope of the quadratic function is always negative. The fourth and fifth figures show the time series of the stock index volatility and the IV volatility respectively. Notice that our approach can be used to provide a forecast of realized short term volatility using as few as three strike prices. In contrast, the well known VIX methodology is not considered to be accurate with only three strike prices. As a result, the VIX methodology is not used for currency options where at most 5 strike prices are liquid at each maturity. The sixth figure shows the time series of the estimated correlation coefficient. The estimated correlation coefficient is below -1 around 25% of the time. This is evidence of either factor model mis-specification or a signal that skew trading would be profitable. Further empirical work on the profitability of such a skew trade is required.

The first table shows the daily cross-sectional linear regression results of regressing the implied variance rate on just a constant. The Black model predicts that this fit would be perfect. The R^2 is in fact zero suggesting the strong need for an adjustment to this Nobel prize winning work. The second table shows that the adjustment to the Black model that we propose leads to R^2 on the order of 95%. While we believe that an alternative specification of risk-neutral IV dynamics can produce an even higher R^2, the existence of an alternative dynamical specification which retains a probabilistic interpretation for the factor arbitrage-free implied variance rate remains open. Panel A in Tables 1 and 2 reports the cumulative distribution of variables of interest across all of the 73,097 cross sectional linear regressions for each day and for each maturity date. The variables including days until expiration, regression coefficients, R^2, and number of observations. Each column is independent, i.e. a 1% cumulative distribution in days until expiration probably does not correspond to a 1% cumulative distribution in σ^2 or a 1% cumulative distribution in the number of observations. Panel B in Tables 1 and 2 reports the median of the variables of interest within each maturity bucket. For example, 1m denotes all options with time to maturity of less than 1 month and 3m denotes all options with time to maturity of less than 3 months, but longer than 1 month.

Finally, Table 3 separates out the instances in which implied correlation is outside $[-1, 1]$ into maturity buckets. The incidence of apparent factor-based arbitrage based on $\rho_t < -1$ clearly increases with time to maturity. If the reason for the slope of ATM IV in k being too negative is the absence of jumps in S, then the maturity pattern should be the opposite of what we observe in Table 2. Nonetheless, the clear pattern in maturity suggests that an IV at some given strike and maturity can be used to anchor an IV at the same strike and a nearby maturity.

Figure 1: **Time-series of σ^2**

This figure plots the time-series of estimated σ^2 using S&P500 index options with 30 to 90 days until expiration. The sample is from 1996 to 2017.

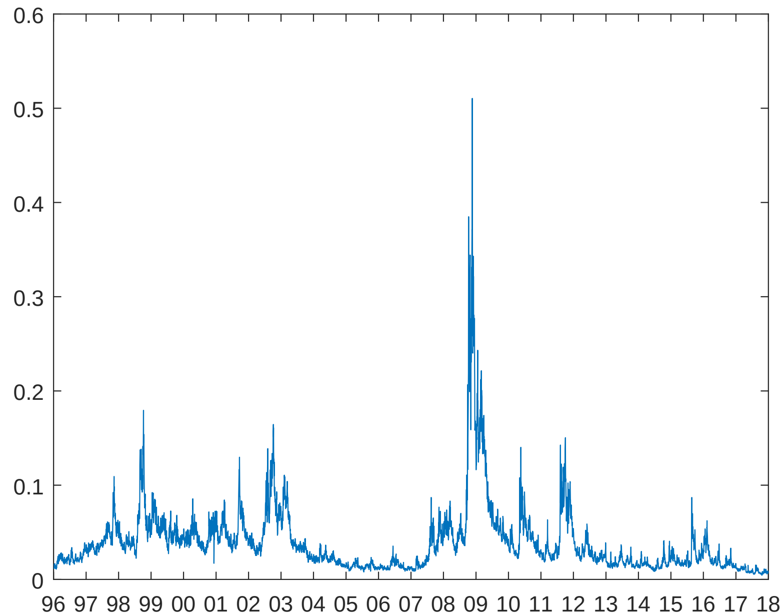

Figure 2: **Time-series of** ω^2

This figure plots the time-series of estimated ω^2 using S&P500 index options with 30 to 90 days until expiration. The sample is from 1996 to 2017.

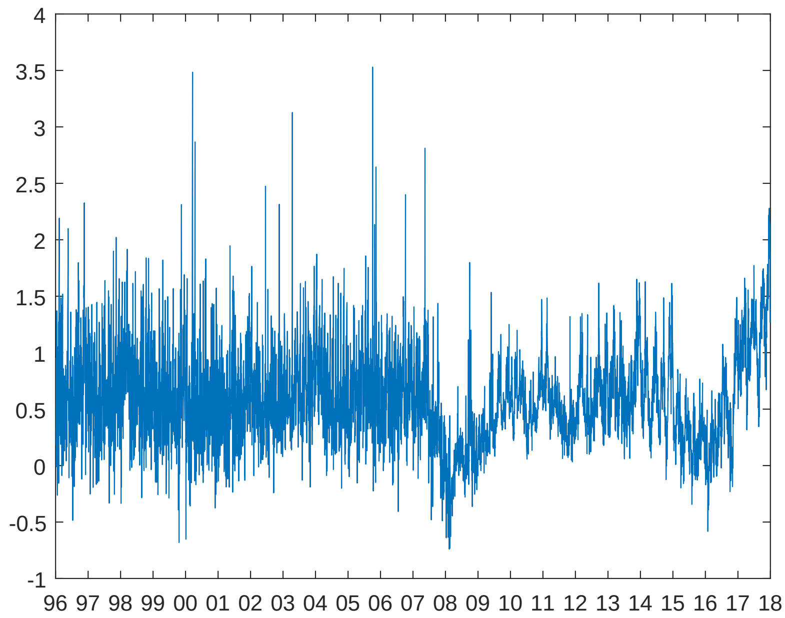

Figure 3: **Time-series of γ**

This figure plots the time-series of estimated γ using S&P500 index options with 30 to 90 days until expiration. The sample is from 1996 to 2017.

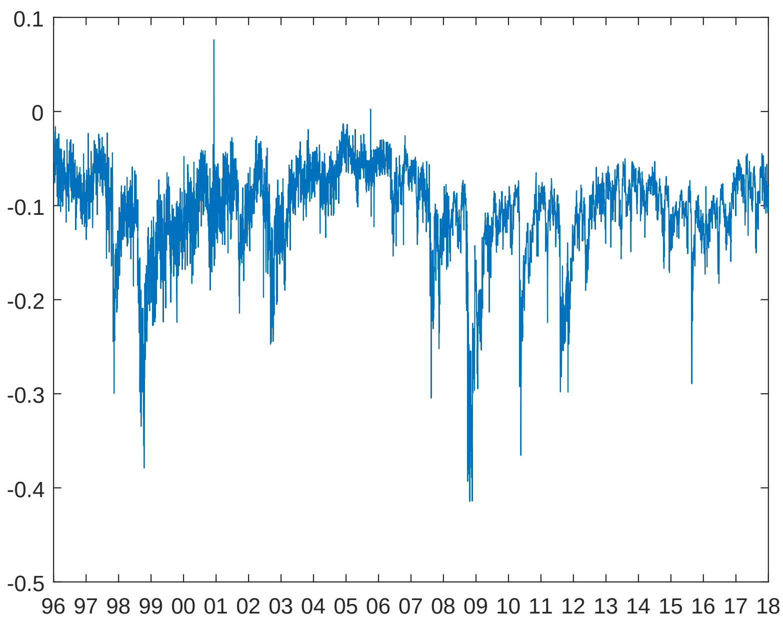

Figure 4: **Time-series of σ**

This figure plots the time-series of estimated σ using S&P500 index options with 30 to 90 days until expiration. The sample is from 1996 to 2017.

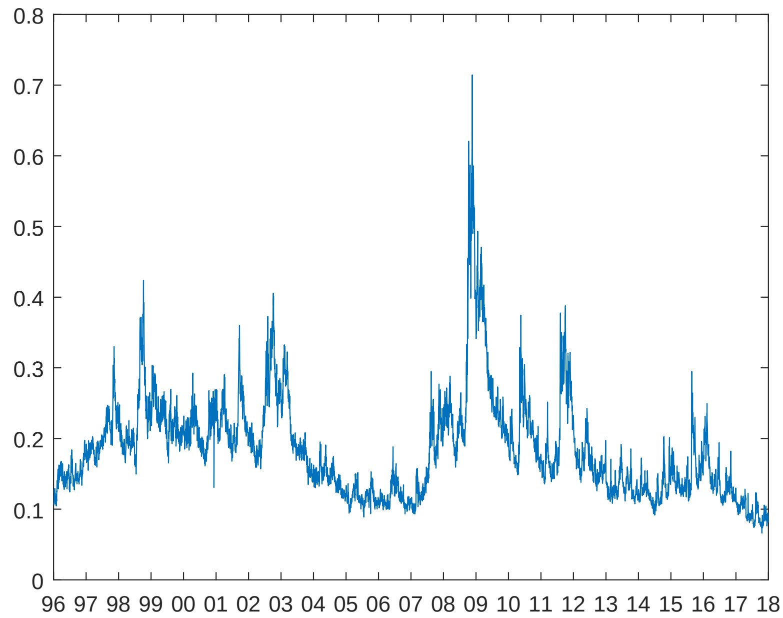

Figure 5: **Time-series of ω**

This figure plots the time-series of estimated ω using S&P500 index options with 30 to 90 days until expiration. Days with negative estimated ω^2 are excluded, which is about 5.6% of the sample. The sample is from 1996 to 2017.

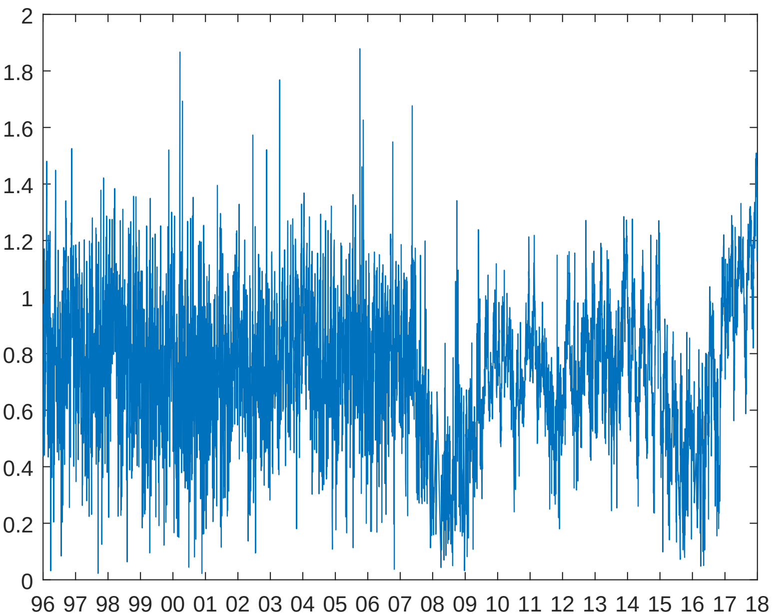

Figure 6: **Time-series of ρ**

This figure plots the time-series of implied ρ, $\rho = \frac{\gamma_t}{\sigma_t \omega_t}$, using S&P500 index options with 30 to 90 days until expiration. The sample is from 1996 to 2017.

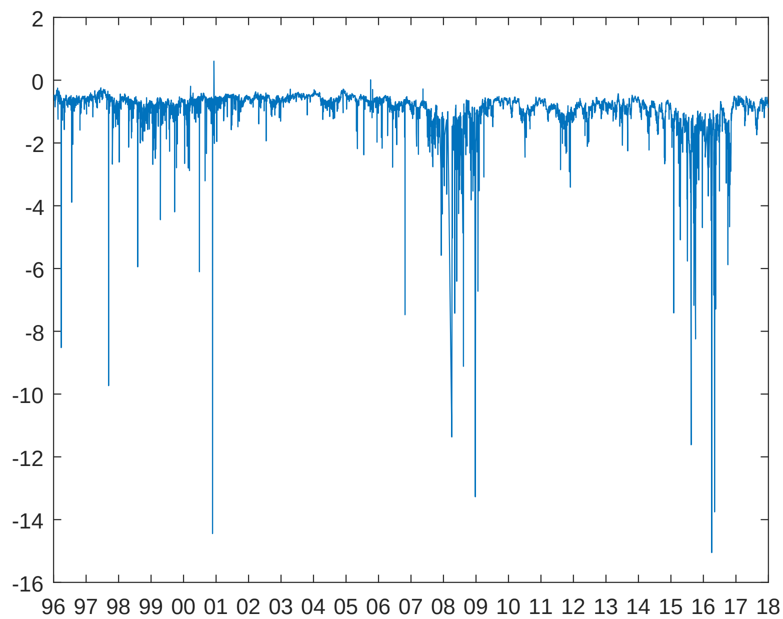

Table 1: **Fitting the Implied Variance Rate with a Constant**

This table shows the daily regressions results of regressing the implied variance rate on a constant. The Black model predicts this fit would be perfect. The regressions are performed using co-terminal S&P500 index options. Panel A shows the distribution of option maturity, coefficients, goodness-of-fit, and number of observations per maturity. Panel B shows the median coefficients, goodness-of-fit, and number of observations each maturity. The sample is from 1996 to 2017.

Panel A: Distribution of Regression Results

	Maturity	Intercept	R^2	Adj R^2	N
1%	9	0.007	0.000	0.000	7
2.5%	11	0.009	0.000	0.000	9
25%	51	0.023	0.000	0.000	17
50%	136	0.036	0.000	0.000	25
75%	319	0.053	0.000	0.000	37
97.5%	785	0.125	0.000	0.000	76
99%	912	0.160	0.000	0.000	88

Panel B: Median Regression Results for Each Maturity

	Maturity	Intercept	R^2	Adj R^2	N
1m	18	0.020	0.000	0.000	26
3m	57	0.028	0.000	0.000	40
6m	128	0.035	0.000	0.000	20
9m	226	0.040	0.000	0.000	22
12m	316	0.042	0.000	0.000	24
2y	525	0.043	0.000	0.000	23
3y	851	0.047	0.000	0.000	36

Table 2: **Regressions of Implied Variance on Moneyness**

This table shows the daily regressions results using specification as in equation (10). The regressions are performed using co-terminal S&P500 index options. Panel A shows the distribution of option maturity, coefficients, goodness-of-fit, and number of observations per maturity. Panel B shows the median coefficients, goodness-of-fit, and number of observations each maturity. The sample is from 1996 to 2017.

Panel A: Distribution of Regression Results

	Maturity (days)	Intercept	k	k^2	R^2	Adj R^2	N
1%	9	0.005	-0.598	-0.504	0.950	0.941	7
2.5%	11	0.007	-0.455	-0.103	0.982	0.979	9
25%	51	0.018	-0.203	0.044	0.998	0.998	17
50%	136	0.029	-0.142	0.132	0.999	0.999	25
75%	319	0.046	-0.099	0.486	1.000	0.999	37
97.5%	785	0.122	-0.048	4.534	1.000	1.000	76
99%	912	0.157	-0.040	7.280	1.000	1.000	88

Panel B: Median Regression Results for Each Maturity

	Maturity (days)	Intercept	k	k^2	R^2	Adj R^2	N
1m	18	0.016	-0.248	1.852	0.998	0.998	26
3m	57	0.022	-0.195	0.477	0.999	0.999	40
6m	128	0.028	-0.157	0.188	0.999	0.999	20
9m	226	0.034	-0.125	0.099	0.999	0.999	22
12m	316	0.035	-0.109	0.064	0.999	0.999	24
2y	525	0.037	-0.084	0.031	0.999	0.999	23
3y	851	0.039	-0.073	0.010	0.999	0.999	36

Table 3: **Frequency of ρ Beyond the Boundary**

This table shows the frequency of the implied ρ larger than 1 or lower than -1, violating the boundary for correlation, for each maturity. The sample is from 1996 to 2017.

	1m	3m	6m	9m	12m	2y	3y
Frequency of $\rho > 1$	0	0	0	0	0	0	0
Frequency of $\rho < -1$	0.145	0.248	0.375	0.498	0.592	0.699	0.914

6 Summary and Further Research

When three or more IV's are known at a given maturity, we developed an IV curve $I_t(K), K > 0$, which has five properties:

1. When exactly 3 IV's are given, $I_t(K), K > 0$ interpolates the three given implied volatilities.
2. For $K > 0$, $k \equiv \ln(K/S)$, and fixed S, $v_t^2(k) \equiv I_t^2(K)$ is a simple explicit algebraic function of k.
3. For $k \in \mathbb{R}$, $v_t^2(k)$ has a simple probabilistic interpretation.
4. $I_t(K), K > 0$ is free of factor-based arbitrage.
5. For $k \in \mathbb{R}$, $v_t^2(k)$ is testable, i.e. it can be rejected using data.

Empirical tests of the restricted quadratic function showed good support for 2 of the 3 restrictions. The third restriction was not always supported since the implied correlation was below negative one about 25% of the time. It is unclear at this time whether the resulting skew trade would be profitable in these instances. Further theoretical analysis and empirical research is required.