

# Implied Volatility And Skewness Surface\*

Bruno Feunou<sup>†</sup>    Jean-Sébastien Fontaine<sup>‡</sup>    Roméo Tédongap<sup>§</sup>  
*Bank of Canada*      *Bank of Canada*      *ESSEC Business School Paris-Singapore*

First Version: November 2008

This Version: October 2013

## Abstract

The Homoscedastic Gamma [HG] model characterizes the distribution of returns by its mean, variance and an independent skewness parameter. The HG model preserves the parsimony and the closed form of the Black-Scholes-Merton [BSM] while introducing the implied volatility [IV] and skewness *surface*. Varying the skewness parameter of the HG model can restore the asymmetry of IV curves. Practitioner's variants of the HG model improve pricing (in-sample and out-of-sample) and hedging performances relative to practitioners' BSM models, with as many or less parameters. The pattern of improvements in Delta-Hedged gains across strike prices accord with predictions from the HG model. These results imply that expanding around the Gaussian density does not offer sufficient flexibility to match the skewness implicit in options. Consistent with the model, we also find that conditioning on implied skewness increases the predictive power of the volatility spread for excess returns.

**Keywords:** SP500 Options, Implied Skewness, Implied Volatility, Volatility Spread, Delta-Hedged Gains.

**JEL Classification:** G12, G13.

---

\*An earlier version of this paper was circulated and presented at various seminars and conferences under the title "The Equity Premium and the Volatility Spread: The Role of Risk-Neutral Skewness". We thank Peter Christoffersen, Redouane Elkamhi, René Garcia, Scott Hendry, Steve Heston, Teodora Paligorova and Jun Yang for their comments. We thank seminar participants at the Bank of Canada, Duke University, CIRANO 2009 Financial Econometrics Conference, Econometric Society 2009 NASM, EFA 2009 and the Fifth International Conference MAF 2012.

<sup>†</sup>Bank of Canada, feun@bankofcanada.ca.

<sup>‡</sup>Bank of Canada, jsfontaine@bankofcanada.ca.

<sup>§</sup>Corresponding Author; ESSEC Business School Paris-Singapore, tedongap@essec.edu.

# I Introduction

The price of an option in the Black-Scholes-Merton [BSM] model can be inverted uniquely for the volatility parameter.<sup>1</sup> The curve describing this implied volatility [IV] across strike prices delivers a transparent comparison of option prices. For traders, this curve describes the time value of different options and provides a direct indication of relative valuations, since it is essentially insensitive to changes in the price of the underlying asset. The shape of the IV curve is largely interpreted in terms of the skewness implicit in option prices.

We study the Homoscedastic Gamma model [HG] in which innovations of market returns are parameterized by their mean, variance and skewness. The skewness parameter can be chosen independently and we nest the Black-Scholes-Merton [BSM] case if skewness is zero. Indeed, the HG model preserves BSM's parsimony and closed-form option prices. We then introduce the implied volatility and skewness *surface*, an extension of the IV curve. Repeating the inversion of option prices for an IV curve across values of skewness delivers the implied volatility and skewness surface. The surface provides a transparent understanding of IV curve variations in term of skewness. There is a smooth relationship between volatility and skewness: negative (positive) skewness increases (decreases) the implied volatility of out-of-the-money calls and decreases (increases) the implied volatility of in-the-money calls. We draw two important conclusions. First, the HG model can restore the symmetry of the observed IV curve and, second, the level of the IV curve also depends on skewness. These conclusions have important implications for pricing and hedging options.

We test the pricing and hedging implications of the simple HG and BSM models as well as their Practitioner's variants [P-HG and P-BSM]. We interpret these variants as expansions around the HG and the Normal distributions, respectively. In addition, we develop restrictions ensuring the identification of the skewness and kurtosis parameters with the true underlying risk-neutral parameters (in the case of the P-HG model). We show that P-HG models in general, and these parameter restrictions specifically, improves out-of-sample pricing and hedging results. Moreover, we show that much of the improvement comes from variations of skewness rather than kurtosis. This imply that density expansion around the Gaussian density, such as the P-BSM model, does not properly capture the skewness present in option data. Another way to view the hedging results is to consider the results of Bates (2005) and Alexander and Nogueira (2005). Essentially, for any contingent claim that is homogenous of degree one, which is the case here, partial derivatives of its price with respect to the change in the underlying (i.e.  $\Delta_t$ ) can be computed, model-free, by taking partial derivatives of observed option prices with respect to strike prices. This implies that models that price options similarly, say if expansion around the Normal and Gamma density are equivalent, then they should have similar hedging implications. The relative hedging performances of the P-HG model delivers a better fit of the true underlying conditional returns distribution but with no increase in implementation costs.

---

<sup>1</sup>The BSM option price formula is a function of the strike price, stock price, interest rate, maturity of the option and of anticipated volatility but only the latter is unobservable.

The HG model also delivers a sharp predictions for Delta-Hedged gains from holding options. The impact of (negative) skewness on the sensitivity of option prices to changes in the underlying index (i.e.  $\Delta_t$ ), relative to the BSM case, is larger for out-of-the-money options than for in-the-money options. Hence, we compare realized gains using  $\Delta_t$  from the BSM and the HG models, respectively. Overall, we find that gains are closer to zero when using the HG model. Moreover, the impact of skewness is highest for short maturity deep out-of-the money options. As predicted, the impact decreases with moneyness and with maturity. The evidence complements Bakshi and Kapadia (2003). They show that stochastic volatility risk explains much of Delta-Hedged gains from holding near-the-money options. Our results suggests that a substantial fraction of Delta-Hedged gains from holding out-of-the-money options can be attributed to non-normality (i.e. jumps) in stock index returns.

Finally, we follow Christoffersen et al. (2009) and provide a Stochastic Discount Factor [SDF] under which stock returns are HG under both the historical and risk-neutral probability measures. This delivers a closed-form analysis of change in the expected returns under both measures (i.e. the equity premium), changes of returns volatility (i.e. the volatility spread) and changes of the skewness of returns. This delivers a sharp prediction about the relationship between the risk premium, the volatility spread and skewness. The equity premium is equal to twice the *ratio* of the volatility spread to skewness. Using volatility and skewness implicit in option prices, we can perform regressions of SP500 excess returns on the ratio of the volatility spread to skewness. We find improvement in predictive power and coefficients have the correct sign and magnitude. In short, theory and evidence shows that the predictive power of the volatility spread analyzed in a recent literature (see e.g. Bollerslev et al. (2008)) is conditional on skewness. Reversing the relationship, and interpreting the volatility spread as the returns on a portfolio of options, we show that a version of the CAPM conditional on skewness “explains” the returns on the volatility spread portfolio. In the light of the question posed in Carr and Wu (2009) regarding which factor may explain the variance premium we argue that skewness plays an important role. Moreover, while different channels can generate skewness and a volatility spread, these have different implications for link, documented here, between the risk premium, the volatility spread and skewness.

### *Related Literature*

The stylized observations that IV curves typically display a smile, a skewed smile or a smirk have been interpreted as evidence of skewness and kurtosis in the underlying risk-neutral distribution of stock price (e.g. Rubinstein and Jackwerth (1998) ). In practice, the importance of skewness for pricing stock index options has been highlighted in the empirical works of Bakshi et al. (1997), Bates (2000) and Christoffersen et al. (2006). However, it is generally difficult to invert option prices and obtain estimates of implied volatility or implied skewness. In most cases, volatility and skewness are not independent or, else, option prices are not available in closed-form, rendering inversion or estimation computationally expensive. Then, although the increased sophistication allows for a better fit of observed IV curves, our understanding of skewness remains incomplete. In particular,

the linkages between skewness, implicit from option prices, the risk premium, measured from equity returns, and the volatility spread remains elusive. The i.i.d. case leads to a stylized model but allows us to maintain parsimony and analytical tractability.

Option pricing based on a Gram-Charlier expansion also offers direct parametrization of skewness and kurtosis (Jarrow and Rudd (1982), Corrado and Su (1996), Potters et al. (1998)). However, approximations of the underlying risk-neutral density often turn negative implying that estimated values of cumulants do not belong to a true distribution. Jondeau and Rockinger (2001) offer a natural remedy and impose a positivity constraint on the estimated density. This is not innocuous. The range of admissible skewness values is restrictive for option pricing applications.<sup>2</sup> Finally, models based on Gram-Charlier do not provide a change of measure linking the historical and risk-neutral measure.<sup>3</sup>

Bakshi and Madan (2000) provide a non-parametric measure of skewness (and other higher-order moments) implicit from option prices. This was exploited by Bakshi et al. (2003), who focus on measures of skewness in the cross-section and on the link with index skewness. Dennis and Mayhew (2000) consider determinants of the cross-section of skewness and Rompolis and Tzavalis (2008) attribute the bias in volatility regressions to the risk-neutral skewness. Christoffersen et al. (2008) explores the information content of option data for future stock betas. However, the pricing or hedging implications of skewness for option prices cannot be handled within this model-free framework.<sup>4</sup>

The relationship between the volatility spread and the equity premium has been attributed to variance risk (Bakshi and Kapadia (2003), Bollerslev et al. (2008), Carr and Wu (2009)) or to a left-skewed and fat-tailed returns distribution (Bakshi and Madan (2006), Polimenis (2006)).<sup>5</sup> While these different channels explain the volatility spread, they do not have the same implications for risk-neutral skewness. Our results suggest that an understanding of the volatility spread, and its relationship with the compensation for risk, demands an understanding of risk-neutral skewness. Intuitively, both the price of risk and the volatility spread are related to the risk-neutral skewness. This should help discriminate across competing theories of the observed volatility spread. Clearly, understanding the source of risk-neutral skewness is a key research objective.

The rest of the paper is organized as follow. Section II introduces the Homoscedastic Gamma model [HG] as well as the SDF and contains the main asset pricing implications. In particular, it contains the mapping between parameters under each measure and derives the option pricing

---

<sup>2</sup>Jondeau and Rockinger (2001) establish that their restriction imply that skewness takes values within  $(-1.0493, 1.0493)$ . León et al. (2006) establishes the impact of this restriction for option pricing.

<sup>3</sup>Note also that closed-form option prices typically result from a first-order approximation. This may not be relevant in practice for option pricing but the impact of this approximation on estimates of implied skewness has not been discussed.

<sup>4</sup>Note, also, that this approach requires approximations of integrals over the moneyness domain. Although Dennis and Mayhew (2000) consider the impact of sampling error under the null of the BSM model, the accuracy of skewness estimates are unknown in the presence of measurement errors or in a non-gaussian setup.

<sup>5</sup>Bakshi and Madan conclude that historical skewness do not play an important role in the determination of the volatility spread but they do not consider risk-neutral skewness.

function. Section III presents the data. Section IV introduces the implied volatility and skewness surface and explores its empirical properties. We introduce practitioner’s analog in Section V and compare their in-sample, out-of-sample and hedging performances in Section VI. Section VII perform tests the implications for the equity premium and the volatility spread, as well as the implications for Delta-Hedged gains, and discusses the results in the context of equilibrium model. Section VIII concludes.

## II The Homoscedastic Gamma Model

This section studies the Homoscedastic Gamma model for stock returns. Our specification is a close analog to Heston (1993) which is also based on the Gamma distribution. The model possesses three crucial properties that makes it a natural choice. First, skewness is parameterized directly and is independent of the mean and variance. Second, its density and characteristic functions are known in closed-form. Third, the distribution of returns remains HG for all investment horizons under both the historical and the risk-neutral probability measures whenever the SDF is exponential in aggregate wealth. In particular, this delivers an explicit mapping between moments under each measures. Finally, we obtain closed-form prices for European options of any maturity as a function of volatility and skewness. We can then efficiently invert option prices to obtain implied volatility and skewness surfaces. Indeed, when setting skewness to zero our model simplifies to the BSM and we recover the usual BSM implied volatility curve.

### A Returns Under the Risk-Neutral Measure

We assume that stock prices,  $S_t$ , follow a discrete-time process whereas the logarithm of gross returns,  $R_t$ , over an interval of time  $\Delta$ , say, follows

$$\begin{aligned} R_{t+\Delta} &\equiv \ln(S_{t+\Delta}/S_t) = \mu^* \Delta + \sqrt{\sigma^{*2}\Delta} \varepsilon_{t+\Delta}^* \\ \varepsilon_{t+\Delta}^* &\sim SG(\alpha^*(\Delta)), \end{aligned} \tag{2}$$

under the risk-neutral measure where  $\mu^*$  and  $\sigma^{*2}$  are the risk-neutral drift and variance, respectively. Return innovations,  $\varepsilon_{t+\Delta}^*$ , follow a Standardized Gamma [SG] distribution with zero mean, unit variance and skewness  $\alpha^*$ . The SG distribution is defined in terms of the Gamma distribution,  $\Gamma(k, \theta)$ , as

$$X \sim SG(\alpha) \Leftrightarrow \frac{2}{\alpha}(X + \frac{2}{\alpha}) \sim \Gamma\left(\frac{4}{\alpha^2}, 1\right), \tag{3}$$

where the scale parameter is fixed to  $\theta = 1$ . Given that the Gamma distribution has mean  $k\theta$ , variance  $k\theta^2$  and skewness  $2/\sqrt{k}$ , it follows that one-period returns in the HG model have mean  $\mu^*\Delta$ , variance  $\sigma^{*2}\Delta$  and skewness  $\alpha^*(\Delta)$ . We express skewness as function of  $\Delta$  to reflect the choice of the interval’s length. A key simplifying assumption is that the conditional distribution of returns does not vary through time. Still, the model could be thought as holding conditionally, with parameters  $\mu_t$ ,  $\sigma_t$  and  $\alpha_t$  indexed by time.

We provide a change of measure for which the historical distribution of stock returns also belongs to the HG family. The result holds when the SDF is exponential-affine in aggregate wealth returns, which is the case in economies with power utility. Under this assumption, we obtain transparent interpretations of risk-neutral moments in terms of the historical moments and of the compensation for risk. This differs from Heston (1993) who considers the case of an SDF from a CRRA economy. Comparing parameters, we find that the risk-neutral volatility is greater than the historical volatility when the equity premium is positive and skewness is negative. Also, the volatility spread increases with the equity premium and with the negative asymmetry of returns. When skewness is zero, and returns are Gaussian, only the mean is shifted and the variance is the same under both measures.

First, assume that aggregate returns follow a HG distribution under the historical measure

$$R_{t+\Delta} \equiv \ln(S_{t+\Delta}/S_t) = \mu \Delta + \sqrt{\sigma^2 \Delta} \varepsilon_{t+\Delta}, \quad (4)$$

where  $\varepsilon_{t+\Delta} \sim SDG(\alpha(\Delta))$ . Next, define the SDF as

$$M_t = \exp(-\nu(\Delta) \varepsilon_t + \Psi(\nu(\Delta))), \quad (5)$$

for some  $\nu$  and where  $\Psi$  is the logarithm of the conditional moment generating function of  $\sqrt{\sigma^2 \Delta} \varepsilon_{t+\Delta}$ . Then, this SDF defines an Equivalent Martingale Measure (EMM), under which the discounted stock price is a martingale, for a unique  $\nu$ , as stated in the following proposition.

**Proposition 1.** *If stock returns follow Equation 4 and if the Stochastic Discount Factor belongs to the class defined by Equation 5 for some  $\nu$ , then, this SDF defines an Equivalent Martingale Measure for discounted stock prices if and only if*

$$\nu(\Delta) = -\frac{2}{\alpha(\Delta) \sqrt{\sigma^2 \Delta}} + \frac{g(\Delta)}{g(\Delta) - 1}, \quad (6)$$

where

$$g(\Delta) = \exp\left(-\frac{(\mu - r)\Delta}{4} \alpha(\Delta)^2 + \frac{\alpha(\Delta) \sqrt{\sigma^2 \Delta}}{2}\right).$$

See the external appendix for all proofs. This is a direct application of results from Christoffersen et al. (2009). Note that the price of risk,  $\nu(\Delta)$ , converges to the usual result,  $(\mu - r)/\sigma^2$ , when skewness tends to zero. Also, this result does not imply that the EMM is itself unique but that only one solution exists within the class defined by Equation 5.

The following Proposition establishes that stock returns are HG under both measures and characterizes the link between parameters of returns dynamics under each measure.

**Proposition 2.** *If stock returns under the risk-neutral measure follow Equation 4 and if the Stochastic Discount Factor is as in Equation 5 for  $\nu$  given in Proposition 1 then stock returns are given*

by Equation 3 and 4 under the risk-neutral and the historical measure, respectively, with  $\varepsilon_t^* = \varepsilon_t - E_{t-1}^Q[\varepsilon_t]$  and where parameters under both measures are linked as

$$\begin{aligned}\sigma^*(\Delta) &= \frac{g(\beta(\Delta)) - 1}{\beta(\Delta)g(\beta(\Delta))} \\ \mu^*(\Delta) &= \mu + 2\frac{\sigma^* - \sigma}{\alpha^*(\Delta)\sqrt{\Delta}} \\ \alpha^*(\Delta) &= \alpha(\Delta)\end{aligned}$$

where we use  $\beta(\Delta) = \alpha(\Delta)\frac{\sqrt{\Delta}}{2}$  to simplify the notation. Note that we have  $\sigma^* \rightarrow \sigma$  and  $\mu^* \rightarrow \mu + \frac{1}{2}\sigma^2$  when  $\alpha - \alpha^* \rightarrow 0$ .

Due to risk-aversion and non-normality in returns, the risk-neutral volatility differs from its historical counterpart at any horizon. The volatility spread depends on the degree of returns asymmetry,  $\alpha(\Delta)$  and the degree of risk aversion through the risk-premium,  $(\mu - r)$ , implicit in  $g(\cdot)$ . Whenever skewness is negative and the equity premium is positive, the risk-neutral volatility is greater than the historical volatility (i.e.  $\sigma^* > \sigma$ ). These results are consistent with Bakshi and Madan (2006) and Polimenis (2006). Finally, because of the specific choice of SDF, the risk neutral skewness is the same as the historical skewness.<sup>6</sup>

To see the relationship between  $\nu$  and skewness, consider a first-order expansion of Equation 6 around  $\alpha(\Delta) = 0$ . For small deviations around the symmetric case we have

$$\nu(\Delta) \approx \frac{\mu - r}{\sigma^2} + \frac{1}{2} + \frac{(\mu - r)^2 + \frac{\sigma^4}{12}}{\sigma^3}\beta(\Delta), \quad (7)$$

Note that  $\nu(\Delta)$  tends toward the usual result,  $\frac{\mu - r}{\sigma^2}$ , when skewness approaches zero. Then, as expected,  $\nu$  can be interpreted as the price of risk. Moreover, it is a function of the equity risk premium, of the volatility and of skewness.

Another way to see the link between the equity premium and the volatility spread is to note that

$$\ln(S_{t+\Delta}/S_t) = \mu\Delta + 2\frac{\sigma^* - \sigma}{\alpha^*\sqrt{\Delta}} + \sqrt{\sigma^{*2}\Delta}\varepsilon_{t+\Delta}^*,$$

where the middle term converges to zero as skewness approaches zero.<sup>7</sup> Taking expectations and re-arranging reveals the following important restriction between the equity premium, the volatility spread and the risk-neutral skewness,

$$E_t^P[\ln(S_{t+\Delta}/S_t)] - E_t^Q[\ln(S_{t+\Delta}/S_t)] = -2\frac{\sigma^* - \sigma}{\alpha^*\sqrt{\Delta}}. \quad (8)$$

---

<sup>6</sup>One can show that an SDF exists such that the returns distribution belongs to the HG family under both measures with *both* the variance and the skewness parameter shifted. However, this SDF is not in general within the exponential-affine class and the link between moments is not transparent.

<sup>7</sup>In the limit, as skewness becomes zero, stock returns follow the usual square-root process.

In the HG model, the volatility spread is solely due to the presence of skewness and not to the presence of a variance risk premium whereas volatility is time-varying and priced by investors. Indeed, the volatility spread *and* the equity premium increase when skewness is more negative. Then, in the presence of non-normality the predictive power of the volatility spread is conditional on skewness. Theory predicts that the correct predictor of the equity premium is the ratio of the volatility spread to skewness instead of the spread itself. In particular, in predictive regressions, the constant should be zero and the predicted value for the coefficients is -2 (see Section VII).

### C Option Prices

We are now ready to provide a closed-form price for European style contingent claims on a stock. This simple homoscedastic model is stable under temporal aggregation. That is, if returns over two successive intervals follow a SDG distribution then returns over the sum of the intervals also follow a SDG distribution. This is a key property to obtain closed-form option prices for all maturities. Consider (log) stock returns over an arbitrary investment horizon  $H$ . Define  $M \equiv \frac{H}{\Delta}$  as the number of time steps over this horizon. Then,

$$\begin{aligned} R_{t,M} &\equiv \sum_{j=1}^M R_{t+j\Delta} = \ln(S_{t+\Delta M}/S_t) \\ &= \mu^* M \Delta + \sigma^* \sqrt{\Delta M} \varepsilon_{t,M}^*, \end{aligned}$$

where the return innovation,  $\varepsilon_{t,M}^*$ , is given by<sup>8</sup>

$$\varepsilon_{t,M}^* \equiv \sum_{j=1}^M \frac{\varepsilon_{t+j\Delta}^*}{\sqrt{M}} \sim SDG(\alpha^*(\Delta)/\sqrt{M}).$$

A no-arbitrage price,  $C_t(K, H)$ , of a European call option with strike price  $K$  and maturity  $H$  can be obtained from the discounted risk-neutral expectation of the terminal payoff,

$$C_t(K, H) = E_t^Q [\exp(-rH) \max(S_{t+H} - K, 0)].$$

As usual, the solution is function of the other model parameters: the risk-free rate,  $r$ , the risk-neutral volatility,  $\sigma^*(\Delta)$ , and the scaled skewness  $\beta(\Delta)$ . Moreover, the solution depends on the direction of asymmetry. Specifically, the case with no skewness corresponds to the BSM formula while we have the following proposition otherwise.

**Proposition 3.** *If the logarithm of gross stock returns follows a Homoscedastic Gamma process under the risk-neutral measure, as in Equation 3, then the price of a European call option is*

$$C_t(K, H) = S_t C_{1,t} - e^{(-rH)} K C_{2,t}, \tag{9}$$

---

<sup>8</sup>This follows directly from the fact that the Gamma distribution is summable.



where, if the skewness is negative (i.e.  $\alpha(\Delta) < 0$ ),

$$C_{1,t} = P\left(\frac{H}{\beta(\Delta)^2}, d_1(\Delta)\right), C_{2,t} = P\left(\frac{H}{\beta(\Delta)^2}, d_2(\Delta)\right), \quad (10)$$

and, if the skewness is positive, (i.e.  $\alpha(\Delta) > 0$ ),

$$C_{1,t} = Q\left(\frac{H}{\beta(\Delta)^2}, d_1(\Delta)\right), C_{2,t} = Q\left(\frac{H}{\beta(\Delta)^2}, d_2(\Delta)\right), \quad (11)$$

The functions  $P(a, z)$  and  $Q(a, z)$  are the regularized gamma functions<sup>9</sup> defined by

$$P(a, z) = \frac{\gamma(a, z)}{\Gamma(a)}, Q(a, z) = \frac{\Gamma(a, z)}{\Gamma(a)},$$

respectively, with  $\gamma(a, z)$  and  $\Gamma(a, z)$  the upper and the lower incomplete gamma functions<sup>10</sup> and where  $d_1$  and  $d_2$  are defined as

$$d_2(\Delta) = \frac{\ln(K/S_t) - \left(r_f + \frac{\ln(1-\beta(\Delta)\sigma^*(\Delta))}{\beta(\Delta)^2}\right) H}{\beta(\Delta)\sigma^*(\Delta)}, d_1(\Delta) = d_2(\Delta)(1 - \beta(\Delta)\sigma^*(\Delta)).$$

### III Data

This section introduces the data and presents some summary statistics. We use prices of call options on the S&P500 index observed on each Wednesday in the period from 1996 to 2004. Using Wednesday observations is common practice in the literature (e.g. Dumas et al. (1998)) to limit the impact of holidays and day-of-the-week effects. Consequently, the return horizon in Equation 3 is set to one week in the following. We exclude observations with less than 2 weeks to maturity, no bid available or with zero transaction volume. We also filter observations for violation of upper and lower pricing bounds on call prices.

Table II displays the number of contracts, the average call price and the average implied volatility across moneyness (Panel (a)), across maturity (Panel (b)), and a detailed cross-tabulation across moneyness and maturity (Panel (c)). The Black-Scholes IV curve is asymmetric in the overall sample, displaying a rising pattern with moneyness, and signaling a sharp left skew in the risk-neutral distribution of returns. Also, the IV curve is flat, or slightly decreasing, with maturity. Disaggregation reveals variations in the shape of the IV curve at different maturities. Starting from the shortest maturity, the IV curve initially follows an asymmetric smile with higher volatility values for in-the-money options. Hereafter, the asymmetry increases as we consider longer maturities and the (average) IV curve eventually becomes monotone in moneyness for the longer maturities.

Note that the composition of the sample varies with maturities. Out-of-the-money contracts

---

<sup>9</sup>We use the standard notation for the regularized gamma functions,  $P(a, z)$  and  $Q(a, z)$ , possibly at the cost of some confusion with the usual notations for the historical and risk-neutral probability measures  $P$  and  $Q$ .

<sup>10</sup>Note that we have  $P(a, z) + Q(a, z) = 1$ , which is a convenient property when computing derivatives (see below).

dominate for long maturities while in-the-money contracts dominate for short maturities. This is due to the issuance pattern of new option contracts. Newly issued, long-maturity call options are typically deep-out-the-money, in anticipation of the index upward drift through time. As we consider shorter maturities, the composition becomes more balanced. At the shortest horizon, most call options are deep in-the-money, since the exchange does not regularly issue short horizon out-of-the-money call options. This implies that the average IV curve reflects, in part, a composition bias with most in-the-money options having short maturities and most out-of-the-money options having long maturities. Because short maturity options have higher implied volatility on average, this makes the average IV curve more smirked.<sup>11</sup>

## IV Implied Volatility and Skewness Surface

In the context of the BSM model, it was recognized early that inverting option prices for the volatility parameter provided a good measure of future returns volatility. However, the HG model offers a separate parametrization for volatility and skewness allowing us to easily measure both the volatility and skewness implicit in option prices.<sup>12</sup> In this section, we study the trade-offs involved between volatility and skewness when fitting option prices. We first analyze how the implied volatility curve varies across different values of skewness and, second, how the implied skewness curve varies with volatility. The results are intuitive. The impact of skewness on implied volatility is asymmetric, depending both on the sign of skewness and of moneyness. In particular, negative skewness tilt a smirked IV curve toward a symmetric smile. On the other hand, the impact of volatility on implied skewness displays a more complex pattern.

An important conclusion from this section is that the HG model exhibits enough flexibility to restore the symmetry of the volatility smile. In other words, variations of the IV curve can be interpreted directly in term of skewness within the HG model. Moreover, both the level and the shape of the IV curve are sensitive to the choice of the skewness parameter. In particular, this implies that empirical studies of the volatility spread, who relies on the level of at-the-money BSM implied volatility, are affected by measurement errors due to the impact of skewness.

### *A Inverting The Implied Volatility and Skewness Surface*

Volatility and skewness cannot be inverted uniquely from a single option price. Instead, for each strike price, the HG model implies a function describing the set of volatility and skewness pairs matching the observed price: a volatility-skewness curve. This is in contrast with the BSM model where any given option price can be inverted uniquely for the volatility parameter. Of course, if the HG model is true, using options with different strike prices would identify uniquely a volatility-skewness couple. In fact, only two different strike prices would be sufficient for this purpose. In

---

<sup>11</sup>This highlights the importance of using a model that can handle maturity differences. In particular, models based on density approximation are not robust to this composition effect.

<sup>12</sup>See Bates (1995) for a review of the literature on the forecasting of volatility using option prices and Andersen et al. (2005) for a review of volatility measurement from stock returns. See Kim and White (2003) for a discussion of the lack of robustness of the usual sample skewness estimator

practice, the HG model extends the BSM model in only one direction, allowing for a skewness parameter. Other deviations from the underlying assumptions cause the volatility-skewness curve to vary across moneyness in such a way that no unique couple can match every observed price. Thus, in the HG model, the counterpart to the IV curve is the implied volatility and skewness *surface*. This surface is the representation of the set of volatility and skewness pairs matching the observed option prices for varying strike prices.

To draw the volatility and skewness surface, we first pick a value of skewness from a grid. Then, each day and for each available strike price, we invert the option price for the volatility parameter and obtain an implied volatility curve. As we vary the value of skewness we obtain different IV curves and, together, they yield an implied volatility and skewness surface. A section of this surface at a given value of skewness is one possible IV curve. Each day, these different IV curves are alternative, and equivalent, representations of the data. Each embodies all the information about the distribution of returns and, in addition, measurement errors due to transaction costs, illiquidity and asynchronous trading. The next section provides the results.

### *B Impact Of Skewness on Implied Volatility Curves*

The average volatility-skewness surface is given in Figure 1 in level (Panel (a)) and in percentage deviations from the benchmark BSM IV curve (Panel (b)). Panel (a) displays the usual smirk in the IV curve when skewness is zero. More interestingly, it shows that the average IV curve is flat for values of skewness around -1.<sup>13</sup> Next, consider the deviations from the BSM curve in Panel (b). The case with skewness equal to zero corresponds to a straight line at zero. As we consider values of skewness away from zero, the IV curve is tilted one way or another depending on the sign of return asymmetry considered. For negative values of skewness, the IV curve is tilted toward positive values of moneyness. Conversely, for positive values of skewness, the IV curve is tilted toward negative values of moneyness. In other words, as we shift probability mass toward the left (right) tail of the return distribution, the implied volatilities required to match observed prices increase (decrease) for out-of-the-money calls and decreases (increases) for in-the-money calls thereby tilting the IV curve back toward a symmetric smile. In the extreme cases, allowing for non-zero skewness can raise or decrease measured implied volatility by more than 15% relative to the BSM case. Clearly, the HG model is sufficiently flexible to capture the skewness implicit in option prices.

### *C Results For Different Option Maturities*

Next, Figures 2 (a)-(e) present implied volatility and skewness surfaces within different maturity groups in percentage deviations from BSM values. Starting with skewness equal to zero, which corresponds to the BSM case, we see the the shape of IV curve varies substantially across maturities. As discussed in section III, the average BSM IV curve is a slightly asymmetric smile for short maturities: implied volatility obtained from in-the-money options is higher than for out-of-the-money options. The smile then gradually disappears as we increase maturity and the IV curve

---

<sup>13</sup>The curve is not strictly flat and this may be due to the impact of kurtosis, or to a composition effect. We discuss these possibilities below.

eventually becomes smirked. For negative values of skewness, and for any maturity, the IV curve is tilted toward a symmetric smile. For short maturities, small negative values of skewness are sufficient to establish a symmetric smile. As we increase maturity, however, more negative values are necessary. Looking at deviations from the case with zero skewness (Figure 2) we see that the impact of a given variation in skewness decreases as we increase maturity.

#### *D Impact Of Volatility On Implied Skewness*

Figures 3 (a)-(f) present implied values for skewness across different values of implied volatility. For at-the-money options, there is no tradeoff between volatility and skewness. However, the impact of volatility on implied skewness is asymmetric and highly nonlinear on both sides of the moneyness spectrum. As the volatility of returns decreases, and the probability mass is closer to the mean, the skewness value required to match observed price increases for out-of-the-money options, implying a higher right-tail, but decreases for in-the-money options, implying a lower left-tail. The reverse is true when we increase the value of volatility. In both cases the impact is not monotonic as we move away from at-the-money. Rather, the pattern follows a sharp V-shape, or inverted V-shape, where changes of volatility have no impact on implied skewness for at-the-money options, the largest impact for intermediate moneyness, and a lower impact for distant moneyness. This is likely an indication of a trade-off between the skewness and the kurtosis in the HG distribution to match observed prices. Finally, the impact of volatility on implied skewness rises with the option maturity.

## V Practitioner’s Models

The previous section shows that the implied volatility and skewness surface can be described as the smooth tilting of the IV curve across values of skewness. However, while the HG model provides enough flexibility to match the skewness present in option data, the IV curve typically remains slightly curved. This is likely due to excess kurtosis. In this section, we propose HG-based option pricing formula that are robust to the presence of excess kurtosis. Intuitively, we consider a one-term expansion of the HG distribution that allows for kurtosis. Note the similarity with the P-BSM model. It is a two-term expansion designed to capture observed skewness and kurtosis deviations from the Gaussian case.

The practitioner’s variants of the BSM model [P-BSM] and of the HG model [P-HG] capture deviations from the Gaussian or HG distributions by modeling volatility as a quadratic function of moneyness. That is, in the P-BSM case, we have

$$\sigma(\xi) = \sigma_0(\alpha, \kappa)(1 + \gamma_1(\alpha, \kappa)\xi + \gamma_2(\alpha, \kappa)\xi^2),$$

and, in the P-HG case, we have

$$\sigma(\xi) = \sigma_0(\kappa)(1 + \gamma_1(\kappa)\xi + \gamma_2(\kappa)\xi^2)$$

where  $\xi$  is moneyness and  $\alpha$  and  $\kappa$  are the skewness and excess kurtosis of the risk-neutral distribution, respectively.

The practitioner’s IV curve smooths through the cross-section of option prices, ignores local idiosyncracies and focuses on the impact of higher-order moments. This approach is pervasive because of its empirical performance and, also, because its parameters (i.e.  $\sigma_0$ ,  $\gamma_1$  and  $\gamma_2$ ) are usually interpreted in terms of the variance, skewness and kurtosis of the true underlying risk-neutral distribution. For these reasons, parameters of the IV curve are commonly estimated without restrictions. In the following, we document that estimates of  $\sigma_0$ ,  $\gamma_1$ , and  $\gamma_2$  vary when we allow for skewness. This contrasts with the usual interpretation of  $\gamma_1$  as a measure of skewness. The remainder of the section provides restrictions on parameters of the IV function such that we can recover direct estimates of  $\alpha$  and  $\kappa$  from option prices.

### A Unconstrained IV Curves

We evaluate empirically the impact of skewness on estimated IV curves. That is, we document how the parameters  $\sigma_0(\cdot)$ ,  $\gamma_1(\cdot)$  and  $\gamma_2(\cdot)$  varies when we vary the skewness of the HG density. To do so, fix the value of  $\alpha$  and estimate the P-HG model at each date. That is, choose values of  $\sigma_0$ ,  $\gamma_1$  and  $\gamma_2$  that minimize squared pricing errors. Next, average the unconstrained estimates through time. Finally, repeat the exercise for different values of skewness and trace the relationships between skewness and estimates of  $\sigma_0$ ,  $\gamma_1$  and  $\gamma_2$ .

Figure 4 presents the results across moneyness categories where, for simplicity we defined

$$\xi = \frac{\ln(S/K)(-r\tau)}{\bar{\sigma}\sqrt{\tau}},$$

Panel (a) presents average estimates of  $\sigma_0$ . For contracts maturing at the next settlement date, at-the-money implied volatility is 20% on average when skewness is zero. When skewness decreases to -3, estimates of at-the-value volatility increase to 23%. Intuitively, shifting some probability mass toward one side of the distribution involves a trade-off for pricing in-the-money versus out-the-money options. For a constant level of skewness, this tension can be reduced by an increase in the level of volatility. A similar pattern occurs at longer maturities, but the impact of skewness gradually decreases. Panel (b) presents the results for the asymmetry parameter. In line with intuition we find that  $\hat{\gamma}_1$  varies linearly with the value of  $\alpha$ : both parameters are measures of the underlying skewness. Finally, Panel 4c shows that  $\hat{\gamma}_2$  also varies substantially with skewness but the relationship is not linear.<sup>14</sup>

The impact of skewness on the IV curve parameter implies that the information on the underlying risk-neutral moments will be shared across unrestricted parameters estimates. Furthermore, the fact

---

<sup>14</sup>This contrasts with the theoretical results of Zhang and Xiang (2005). They argue that in the Gaussian case and up to a first-order approximation  $\sigma_0(\beta, \kappa)$  is linear in the risk-neutral volatility,  $\gamma_1(\beta, \kappa)$  is linear in skewness, and  $\gamma_2(\beta, \kappa)$  is linear in kurtosis. However, they assume that the skewness and excess kurtosis of the underlying distribution can be chosen independently while in fact there is a tight link between the two for any given correctly specified density. Moreover, they linearize around the case where  $\sigma = 0$  and this may lead to a poor approximation.

that estimates of  $\alpha$  and of  $\gamma_1$  are (linearly) correlated suggests that they are poorly identified. The following section introduces a framework which lead to restrictions on  $\sigma_0$ ,  $\gamma_1$  and  $\gamma_2$  such that only  $\hat{\alpha}$  can capture the risk-neutral skewness. Absent these restrictions, each parameter estimates of the IV curve,  $\sigma_0$ ,  $\gamma_1$  and  $\gamma_2$ , is affected by skewness and kurtosis (i.e.  $\alpha$  and  $\kappa$ ). In contrast, the restrictions developed below leads to the unambiguous identification of skewness.<sup>15</sup>

### B HG Model With Excess Kurtosis

We now provide a rigorous justification of the P-HG model when the true distribution displays excess kurtosis. We can characterize sufficient restrictions on the parameters of the IV curve such that  $\hat{\alpha}$  is identified as the risk-neutral skewness in this more general model as well. In this context, parameters of the IV curve are restricted to (known) functions of excess kurtosis. In other words, any deviation from a flat IV curve can only be linked to deviations of  $\kappa$  from zero. As a by-product, we obtain an estimator of the kurtosis in excess of the Gamma distribution.

Intuitively, we assume that the true density of returns can be represented by an Edgeworth expansion around the Gamma distribution. This is similar to earlier work using the Gaussian distribution (Jarrow and Rudd (1982), Corrado and Su (1996)) but the Gamma distribution allows an exact match of the first three moments. We then impose the equality of the option pricing formula under the true model and the P-HG model for at-the-money options.

Suppose that the true evolution of stock returns under the risk neutral measure can be described as

$$R_T = (r - \delta^*)T + \sigma^* \sqrt{T}y,$$

where  $\delta^*$  is a risk-adjustment term,  $y$  is a random variable with mean zero, unit variance, skewness,  $\alpha^*$  and kurtosis,  $\lambda_2^*$ . We allow for deviations from HG case and assume that the probability density of  $y$  is given by

$$f(y) = h(y) + \frac{\lambda_2^* - \frac{3\alpha^{*2}}{\sqrt{T}}}{4!} \frac{d^4 h(y)}{dy^4}, \quad (12)$$

where  $h(y)$  is the standardized gamma density. This is a one-term Edgeworth expansion of standardized gamma distribution around the case with no excess kurtosis. If  $y$  is normally distributed, then  $\alpha^* = 0$  and  $\delta^* = \frac{\sigma^{*2}}{2}$ . This approach captures in a rigorous way deviations from a flat skewness and volatility surface documented above. These deviations are linked to fat tails in the distribution of returns in excess of the Gamma distribution but ignores deviations beyond the fourth moment. Our objective here is to derive explicitly the function  $\sigma_0(\kappa)$ ,  $\gamma_1(\kappa)$  and  $\gamma_2(\kappa)$ . Proposition 4 builds on a no-arbitrage argument and provides a closed-form characterization of option prices and of the risk-adjustment term.

**Proposition 4.** *If the logarithm of gross stock returns has the density given by Equation 12, then*

---

<sup>15</sup>Note that merely imposing  $\gamma_1(\alpha, \kappa) = 0$  does not identify an estimator of  $\alpha$  with skewness.

the price of a call option,  $C^*(K, T)$ , with maturity  $T$ , underlying price  $S_0$  and strike price  $K$  is

$$\begin{aligned} C^*(K, T) &= S_0 P(a^*, d_1^*) - e^{-rT} (1 + T^2 \sigma^{*4} \kappa_4) K P(a^*, d_2^*) \\ &\quad + \kappa e^{-rT} K \frac{T^2 \sigma^*}{\frac{1}{2} \alpha^{*3}} \left[ -h''(d_2^*) + \sigma^* \frac{1}{2} \alpha^* h'(d_2^*) - \sigma^{*2} \frac{1}{2} \alpha^{*2} h(d_2^*) \right], \end{aligned}$$

when  $\alpha < 0$  and

$$\begin{aligned} C^*(K, T) &= S_0 Q(a^*, d_1^*) - e^{-rT} (1 + T^2 \sigma^{*4} \kappa_4) K Q(a^*, d_2^*) \\ &\quad - \kappa e^{-rT} K \frac{T^2 \sigma^*}{\frac{1}{2} \alpha^{*3}} \left[ -h''(d_2^*) + \sigma^* \frac{1}{2} \alpha^* h'(d_2^*) - \sigma^{*2} \frac{1}{2} \alpha^{*2} h(d_2^*) \right], \end{aligned}$$

when  $\alpha > 0$ . We define the excess kurtosis,  $\kappa = \frac{\lambda_2 - \frac{6\beta^{*2}}{T}}{4!}$ , and

$$\begin{aligned} d_2^* &= \frac{\ln(K/S_0) - rT - a^* \ln(1 - \sigma \frac{1}{2} \alpha) + \ln(1 + T^2 \sigma^4 \kappa)}{\sigma \beta} \\ d_1^* &= d_2 \left( 1 - \sigma \frac{1}{2} \alpha \right), \quad a^* = \frac{T}{\frac{1}{2} \alpha^2}, \end{aligned}$$

where  $h$  is the density of the standard gamma distribution.

### C Identified practitioner's HG

We are now looking for the restrictions on the parameters of the P-HG model such that estimation of  $\beta$  delivers a convergent estimate of the risk-neutral skewness  $\beta^*$ . Zhang and Xiang (2005) provide the restriction for the case where the Gaussian density is used in the approximation. To find the link between the parameters of the P-HG model with parameters of the true distribution, we impose the following restrictions

$$C^*(K, T) = C(K, T), \quad \frac{\partial C^*(K, T)}{\partial K} = \frac{\partial C(K, T)}{\partial K}, \quad \frac{\partial^2 C^*(K, T)}{\partial K^2} = \frac{\partial^2 C(K, T)}{\partial K^2}$$

when evaluated at-the-money (i.e.  $K = S_0 e^{rT}$ ). These restrictions are given in the external appendix but note that they are trivially satisfied whenever  $\kappa = 0$  since in this case the HG model is true and the IV curve is flat. Of course this corresponds to the case  $\sigma_0 = \sigma^*$ ,  $\alpha = \alpha^*$  and  $\gamma_1 = \gamma_2 = 0$ . We linearize the restrictions around this point (i.e.  $\kappa = 0$ ) and obtain

$$\frac{\sigma_0 - \sigma}{\sigma} = A_1(\sigma, \alpha) \kappa \tag{13}$$

$$\gamma_1 = B_1(\sigma, \alpha) \frac{\sigma_0 - \sigma}{\sigma} + B_2(\sigma, \alpha) \kappa \tag{14}$$

$$\gamma_2 = C_1(\sigma, \alpha) \frac{\sigma_0 - \sigma}{\sigma} + C_2(\sigma, \alpha) \gamma_1 + C_3(\sigma, \alpha) \kappa, \tag{15}$$

where the coefficients are given in the external appendix.<sup>16</sup> Then, small deviations of the underlying density from a HG distribution lead to deviations from a flat implied volatility and skewness surface. This highlights the impact of excess kurtosis on the estimates of  $\sigma_0$ ,  $\gamma_1$  and  $\gamma_2$ . It also makes clear that deviations from a flat IV curve are only due to excess kurtosis. More importantly, these restrictions ensure that  $\alpha$  corresponds to the risk-neutral skewness and that the practitioner’s HG model conforms to the true returns density.

## VI Option Pricing Results

In this section, we estimate each model and compare their performance. The results show that the HG framework substantially improves in-sample, hedging and out-of-sample performances. The improvements are robust if we impose identification of the skewness parameters, as discussed in the previous section. Indeed, the improvements remain when the only deviation from the HG model is an adjustment to kurtosis that is constant through time. Out-of-sample, imposing the identifying restrictions does not degrade pricing performance. In other words, a fixed implied volatility and skewness surface combined with variations in skewness delivers most of the in-sample and out-of-sample improvements.

Overall, our approach delivers a reliable measure of skewness while offering improved forecasting and hedging performance. In contrast, the P-BSM model does not allow for sufficient flexibility to match the skewness implicit in the data and offers lower hedging and out-of-sample performance. While the more general models that allow for variations in excess kurtosis perform better in-sample, these improvements disappear out-of-sample. This implies that skewness captures most of the persistent deviations from the Gaussian case and that excess kurtosis and other deviations are transitory.

### A Description Of Models

We evaluate the basic HG model and the usual P-BSM model. We also include three different versions of the P-HG model based on the quadratic IV curve,

$$\sigma_t(\xi) = \sigma_0(1 + \gamma_1\xi + \gamma_2\xi^2).$$

where the first version, P-HG1, imposes the simple restriction that  $\gamma_1 = 0$ . This is another way to see that the usual interpretation of  $\gamma_1$  as a measure of skewness, while intuitive, is misleading. The second model, P-HG2, imposes the restrictions derived in the previous section and delivers an estimate of skewness robust to excess kurtosis. Finally, P-HG3 is unrestricted.

We also introduce “smoothed” versions of these models where some parameters of the IV curves are held constant through the sample. First, the smoothed version of the P-HG1 model, labeled SP-HG1, still imposes that  $\gamma_1$  is zero but holds  $\gamma_2$  constant through time. Next, SP-HG2 still allows

---

<sup>16</sup>We differ from Zhang and Xiang (2005) who linearize the restrictions around  $\sigma = 0$ . Arguably, linearizing around the HG distribution is likely to provide a better approximation than linearizing around the deterministic case.



Table I: **Summary of Models**

The specification of each model described in the text can be characterized through its parametrization of  $\sigma(\xi)$ ,  $\alpha(\xi)$  and  $\kappa(\xi)$ . The last column gives the number of time-varying parameters in each model.

Model	$\sigma(\xi)$	$\alpha(\xi)$	$\kappa(\xi)$	# Par.
BSM	$\sigma(\xi) = \sigma_{0,t}$	$\alpha = 0$	$\kappa(\xi) = 0$	1
P-BSM	$\sigma(\xi) = \sigma_{0,t}(1 + \gamma_{1,t}\xi + \gamma_{2,t}\xi^2)$	$\alpha(\xi) = 0$	$\kappa(\xi) = 0$	3
HG	$\sigma(\xi) = \sigma_{0,t}$	$\alpha(\xi) = \alpha_t$	$\kappa(\xi) = 0$	2
P-HG1	$\sigma(\xi) = \sigma_{0,t}(1 + \gamma_{2,t}\xi^2)$	$\alpha(\xi) = \alpha_t$	$\kappa(\xi) = 0$	3
P-HG2	$\sigma(\xi) = \sigma_0(\alpha_t, \kappa_t)(1 + \gamma_1(\alpha_t, \kappa_t)\xi + \gamma_2(\alpha_t, \kappa_t)\xi^2)$	$\alpha(\xi) = \alpha_t$	$\kappa(\xi) = \kappa_t$	3
P-HG3	$\sigma(\xi) = \sigma_{0,t}(1 + \gamma_{1,t}\xi + \gamma_{2,t}\xi^2)$	$\alpha(\xi) = \alpha_t$	$\kappa(\xi) = 0$	4
SP-HG1	$\sigma(\xi) = \sigma_{0,t}(1 + \gamma_2\xi^2)$	$\alpha(\xi) = \alpha_t$	$\kappa(\xi) = 0$	2
SP-HG2	$\sigma(\xi) = \sigma_0(\alpha_t, \kappa)(1 + \gamma_1(\alpha_t, \kappa)\xi + \gamma_2(\alpha_t, \kappa)\xi^2)$	$\alpha(\xi) = \alpha_t$	$\kappa(\xi) = \kappa$	2
SP-HG3	$\sigma(\xi) = \sigma_{0,t}(1 + (\gamma_{10} + \gamma_{11}\alpha_t)\xi + (\gamma_{20} + \gamma_{21}\alpha_t)\xi^2)$	$\alpha(\xi) = \alpha_t$	$\kappa(\xi) = 0$	2

for a flexible fit of skewness through time but keep excess kurtosis constant through time. We include this model as a simple way to evaluate the relative importance of skewness and kurtosis for option pricing and hedging. Finally, the SP-HG3 model imposes the following structure on the IV curve,

$$\sigma(\xi) = \sigma_0(1 + (\gamma_{10} + \gamma_{11}\alpha)\xi + (\gamma_{20} + \gamma_{21}\alpha)\xi^2).$$

which is a simple attempt to implement the observation made in Section V that parameters of the IV curve vary with skewness. Finally, estimation is performed through minimization of squared pricing errors in the weekly sample. The following table summarize the various specification.

## B In-Sample RMSE

### B.1 HG And BSM Models

Table III presents in-sample Root Mean Squared Errors [RMSE] where each results is expressed as a percentage of the BSM's RMSE. Panel (a) presents results across moneyness while Panel (b) presents results across maturities. Although the most flexible (i.e. P-HG3) model achieves an RMSE which is 14% of the benchmark, most of the improvement comes from using the HG distribution: the simpler HG model's RMSE is 37% of the BSM's RMSE but with only more parameter measuring skewness.

### B.2 Practitioner's Variants

Interestingly, even with one extra parameter, the P-BSM does not offer much improvement (35% vs 37%) over the straightforward HG model. The models offer similar results across maturities but their performances differ across strike prices. The P-BSM improves pricing for in-the-money options at the expense of larger errors for other moneyness groups. On the other hand, the P-HG1 and the P-HG2 models achieve RMSEs that are 28% and 23%, respectively, but with the same number of

parameters as the P-BSM model. However, in contrast with the P-BSM model, the lower errors for out-of-the-money options are not compensated by higher errors for options that are nearer the money. Thus, models based on the HG distribution appear to offer more flexibility than the practitioner’s BSM in choosing risk-neutral skewness and kurtosis but with equal or less parameters.

Although the naive  $\gamma_1 = 0$  restriction seems reasonable, it fails in practice with larger RMSE. Comparing models, we see that imposing the correct identification constraints (P-HG2) provides substantial improvement over the P-HG1, especially for short maturity, out-of-the-money call options. Finally, with one more parameter, the P-H3 offers much lower in-sample RSME (14%) than any other model across all moneyness and maturity categories.

### B.3 Smoothed Coefficients

Smoothed models have less parameters but the SP-HG2 model still improves (31%) over the P-BSM model but with less parameters. This model has the flexibility to fix skewness from date to date but imposes a constant excess kurtosis. That is, deviations of the IV curve from the HG case are kept constant. Thus, in-sample, a flexible fit of the underlying risk-neutral skewness is key while variations in kurtosis are less important. Finally, while more flexible HG-based models improve the in-sample fit, the next section show that this result is not robust out-of-sample, indicating a relatively minor role for information beyond the third moment.

#### C Out-of-sample RMSE

The improved performance of models based on the HG distribution may be due to over-fitting and may not hold out-of-sample. This section compares the out-of-sample performance of each model. First, we estimate each model from options in a given week.<sup>17</sup> We then fix these parameters and price options observed in the following week. Table IV presents one-week out-of-sample RMSE for each model across strike prices (Panel (a)) and across maturities (Panel (b)).

Out-of-sample, the improvement in fit relative to the BSM decreases for all models. This indicates that some of the deviations from the Gaussian case are transitory. The lowest relative RMSE is now 57%, obtained for the P-HG3 model, with 4 parameters. On the other hand, the worst result is 68%, obtained for the P-BSM model, with 3 parameters. This add to the evidence that the practitioner’s version of the BSM model does not properly fit the persistent skewness and kurtosis present in the data. Strikingly, the SP-HG2 model, which uses 2 parameters and fixes excess kurtosis through the sample, actually improves out-of-sample RMSE (64%) over the more flexible P-HG2 and P-BSM models. Some of the variations in excess kurtosis required to match (in-sample) option prices in this category are transitory, degrading out-of-sample performances. Restricting parameters of the IV curve to capture that part of its variations due to skewness improves the out-of-sample fit.

---

<sup>17</sup>For smoothed model, we estimate parameters that are held constant through the sample in a first pass.

Hedging errors implied by each model may convey more economic significance to risk-managers. Below, we verify that allowing for skewness significantly alter hedging strategy theoretically, and improves hedging results empirically. Also, we verify that any improved hedging performance persists at horizons beyond one week. The SP-HG2 model, with 2 parameters and where skewness is separately identified, offers the next to best performance. This highlights, again, the value of theoretically sound restrictions. Again, we find that the unrestricted P-HG3 model performs best.

### **D.1 Comparing The Greeks**

As in the BSM model, we can compute explicitly the sensitivity of option prices to changes in the underlying parameters, including the sensitivity to changes in skewness. We provide these in the external appendix. These derivatives depend on the direction of asymmetry and everywhere the symmetric case (i.e.  $\alpha = 0$ ) leads to the standard results from BSM. To see the impact of skewness, we draw options sensitivities across strike prices for different values of skewness. In the computations, we use the average values of volatility, of the interest rate and of the index level. Figure 5 presents results for the first and second derivatives with respect to the underlying, Delta and Gamma, as well as the derivative with respect to volatility, Vega. The results are reported in levels in the top panels (Panel (a) to (c)) and in percentage deviations from the symmetric case in the bottom panels (Panel (d) to (f)).

First, the pattern of Delta across moneyness is familiar. The sensitivity is small for deep out-of-the-money options but grows to close to one for deep in-the-money options. Varying skewness does not alter this picture but looking at levels hides significant deviations. At skewness equal to -2.5, which occurs in our sample, short positions in the stock are as much as 20% higher for some out-of-the money or near to at-the money options. Next, the impact on Gamma is dramatic. In the symmetric case, Gamma appears quadratic in moneyness with highest values for at-the-money options. Decreasing skewness lowers Gamma for in-the-money options but increases Gamma for out-of-the-money options. When skewness is -2.5, Gamma is as much as 50% lower then when skewness is zero for in-the-money options and 50% higher for out-of-the-money options. Finally, skewness has an asymmetric impact on the sensitivity of options to variations in volatility. When skewness is -2.5, Vega decreases by more than 20% for out-of-the-money options but increases by nearly 20% for in-the-money options. Clearly, ignoring the impact of skewness can lead to large hedging errors, which is confirmed empirically in the next section.

### **D.2 Comparing Hedging Performance**

We follow Dumas et al. (1998) and compute hedging errors as

$$\epsilon_t = \Delta C_{t,t+h}^{actual} - \Delta C_{t,t+h}^{model}$$

which is a measure of the impact of changes in model errors from  $t$  to  $t+h$  on the hedging strategy.<sup>18</sup> By this measure, a good model delivers hedging errors that are close to zero on average. Table V and Table VI present the results for hedging horizons from one to four weeks ahead (i.e.  $h = 1, 2, 3, 4$ ).

Consider hedging errors at the 1-week horizon (Table Va). First, the BSM model appears to perform well, with hedging errors averaging 1.6 cents. But this hides important disparities across maturities. Average hedging errors range from 36.7 cents for out-of-the-money options to -39 cents for in-the-money options. Moreover, the more flexible P-BSM model has higher overall hedging errors (-4.6 cents) with substantial average errors (-18.8 cents) for the lowest strike prices.

When considering the overall mean and the dispersion of hedging errors across maturities, the best performing models are variants of the P-HG model. Identification restrictions for skewness perform well. In particular, the SP-HG2 model offers both low overall hedging errors and low dispersion across moneyness. Averages remain below 10 cents across strike prices. Table Vb draws a similar picture at the 2-week horizon. The P-BSM model sees its average performance deteriorate to -8.2 cents and mean hedging errors now range from -21.8 to 7.1 cents. Again, HG-based models offer better performance. The SP-HG2 model still offers the best performance: the mean pricing error is 0.002 cents in the entire sample and ranges from -13.6 cents to 8.6 cents across moneyness. Finally, results at the 3 and 4-week horizons (Tables (a) and (b)) quickly deteriorate for the BSM and the P-BSM models. However, the SP-HG2 model still performs well. The overall averages at 3-week and 4-week horizons are -4.3 cents and -2.1 cents.

## *E Discussion*

Overall, the results favor the more general P-HG3 model. It offers lower in-sample and out-of-sample RMSEs as well as better hedging performances at all horizons. This contrasts with the frequent observation that the P-BSM model offers sufficient flexibility. Indeed, option prices based on the HG distribution offer better performance than the P-BSM with as many parameters (P-HG1 and P-HG2) or less (SP-HG2). If we interpret the practitioner's models as expansions around the Gaussian or the Homoscedastic Gamma distributions, the results imply that expanding around the Gaussian density is restrictive and does not offer sufficient flexibility to match the skewness and kurtosis implicit in the data. Moreover, when we consider the sequence of models, we see that imposing restrictions such that skewness is correctly measured and excess kurtosis constant does preserve most of the performance improvement.

Another way to view these results is to consider the results of Bates (2005) and Alexander and Nogueira (2005). Essentially, they show that for any contingent claim that is homogenous of degree one, all partial derivatives with respect to the underlying can be computed by taking partial derivatives of option prices with respect to strike prices. This implies that, if the number of observed option prices is arbitrarily large, we can compute delta and gamma exactly from non-parametric derivatives. In practice, however, some parametric model is fitted to observed prices from which

---

<sup>18</sup>This abstracts from the hedging errors due to discrete adjustments. See Galai (1983) for details.

derivatives can be imputed. The hedging performances of the P-BSM and the P-HG models imply that the latter offer a better fit of the true option price curve across the strike continuum and, therefore, a better fit of the true option’s delta and gamma. In other words, the relatively poor fit of skewness by Gaussian-based expansions translates in inaccurate option sensitivity measures and larger hedging errors relative to approximations based on the Gamma density.

The performance of the SP-HG2 model implies that the parametric measure of risk-neutral skewness is relevant. This provides a measure of skewness that is easy to compute and requires less data than a non-parametric measure. Moreover, together with the regression results from Section VII, the importance of skewness for hedging and out-of-sample pricing confirms the key link between the risk premium and volatility shift across moneyness and skewness. Indeed, imposing the additional restriction that excess kurtosis is constant yields the next to best out-of-sample and hedging performances. Interestingly, the estimate of  $\kappa$  is negative (-0.042). Then relaxing the link between kurtosis and skewness allows for more asymmetry to be applied to the data than the benchmark HG model does. This adjustment is significant: to keep kurtosis constant but with  $\kappa$  equal to zero, skewness would have to be reduced (closer to zero) by 0.21. Taken together, the results lead us to adopt the SP-HG2 as our preferred model to measure the option-implied skewness.

## VII Skewness And The Compensation For Risk

### A Implications For The Volatility Spread

When the representative SDF can be approximated by the exponential-shift given in Equation 5 we have a tight link between the price of risk, the volatility spread and skewness. After some manipulation of Equation 8, we obtain

$$\ln(S_{t+i}/S_t) - r_t^{(i)} - \omega_t^{(i)} = -2 \frac{\sigma_t^{*(i)} - \sigma_t^{(i)}}{\alpha_t^{(i)}} + \sigma_t^{*(i)} \varepsilon_{t+i}^*$$

for an investment horizon  $i$  and where  $r_t^{(i)}$  is the risk-free rate for that horizon and  $\omega_t$  is the Jensen adjustment term.<sup>19</sup> In the following, we test this implication of the HG model and its ability to capture the volatility spread and the equity premium. We perform monthly regressions of SP500 (log) excess returns on the ratio of the volatility spread to skewness. The key predictions are that the constant should be zero and that the coefficient should be -2. Each month, Risk neutral volatility and skewness measures are obtained through minimization of squared pricing errors of the HG model. Finally, we measure the historical volatility using the observed monthly realized volatility.

---

<sup>19</sup>This term is a function of both skewness and volatility but the first term of its Taylor expansion is the usual correction in the Gaussian case,  $\frac{1}{2}\sigma^2$ .

## B Regression Results

Table IX presents the results from regressions of excess returns at horizons of 1, 3, 6, 12, 24 and 36 months on the ratio of the volatility spread to skewness.<sup>20</sup> The results are striking. Point estimates for the slope coefficient are close to -2 as predicted by the model. Moreover, at horizons of 3, 6, and 12 months, where we would expect the forward-looking nature of the option-implied estimate to be the most relevant, estimates are -2.24, -2.04 and -2.13, respectively. In fact, at any horizon, we cannot reject the null hypothesis that the coefficient is equal to -2. Next, the constant is not significantly different from zero so that the two most important implications of the model cannot be rejected empirically. Finally, the predictability of excess returns is low at the 1-month horizon (i.e.  $R^2$  is 1.85%) but rises steadily with the horizon, reaching 5.6%, 9.7% and 11.3% at horizons of 6, 12 and 36 months.

For comparison with results available in the existing literature, we also consider regressions on the volatility spread which displays some predictive power at horizons of 9 and 12 months. However, coefficients are not significant at other horizons. Finally, we ask if the volatility spread contains information beyond that revealed by the ratio of the volatility spread and skewness. The results from the regressions are presented in Table IX. Since volatility and the ratio of the volatility spread to skewness are correlated, the coefficients become unreliable, even changing signs. However, their combined predictive power does not rise above that of the volatility to skewness ratio, further supporting the implications of the model.

## C Skewness And Delta-Hedged Gains

Volatility risk, as well the presence of skewness or excess kurtosis, induces a risk premium in index option markets (see e.g. Bakshi and Kapadia (2003)). Holding a call option and hedging its exposure to the underlying with a short position in the index consistently yields losses. For reason of convenience, these results are typically based on explicit hedging strategies derived from the BSM model (i.e.  $\Delta_t$ ). These derivatives are also available in closed-form in the case of the HG model but they also depend on the measured skewness. In fact, to the extent that this risk premium is attributable not to the presence of volatility risk but to presence of skewness and kurtosis, Delta-Hedged gains should be reduced substantially when using the HG-based  $\Delta_t$ . This is relevant because the HG model provides explicit solution for  $\Delta_t$  and can be easily implemented by practitioners.

This sections explores whether the skewness allowed in the simple HG model captures part of Delta-Hedged gains. We perform the following simple exercise. First, purchase a call option and then establish a short position in the index to hedge against the sensitivity to the underlying. Then, hold the position for a week, close all positions and compute gains. We assume that one can borrow

---

<sup>20</sup>Precisely, our measures of risk-neutral moments pertain only to the distribution of returns at a horizons of 12 months or less. Nonetheless, if these moments exhibit persistence, their predictive power will extend to longer horizons as is indeed the case

and invest at the risk-free rate,  $r_t$ . The gains are given by

$$\pi_{t+1} = C_{t+1} - C_t - \Delta_t(S_{t+1} - S_t) - r_t(C_t - \Delta_t S_t) \frac{1}{52} \quad (16)$$

where the hedge ratio,  $\Delta_t$ , is computed from the BSM or from the HG model, respectively.

At each date, we compute the gain for each available option whenever that same option is observed in the following week. We average gains within each maturity and moneyness category and Table VIII provides results from the BSM model (Panel (a)) and from the HG model (Panel (a)). Consistent with previous literature, we see average Delta-Hedged gains from the BSM model are typically negative. Gains are particularly low for out-of-the-money short maturity options but increase with moneyness and maturity. In fact, gains are close to zero for in-the-money options. Using  $\Delta$  based on the HG model reduces Delta-Hedged losses substantially. This shows that skewness captures important information about the compensation for risk offered to option sellers. Results using HG-based  $\Delta_t$  suggests a much lower risk premium (i.e. lower losses from holding calls) for out-of-the-money option for each maturity category while gains are similar to the results from the BSM case for in-the-money options. The results are consistent with the predictions of the model. Recall Figure 5d, which gives the impact of skewness on  $\Delta_t$  across moneyness for different values of skewness relative to the  $\Delta_t$  computed from the BSM model. Clearly, the impact of a negative skewness is largest for out-of-the-money options but small or zero for in-the-money options.<sup>21</sup>

#### D Discussion

We can interpret the results in the broader context of a general equilibrium model. There, the price of risk is determined by preference parameters. In particular, in an economy with power utility,  $\nu$  corresponds to the risk-aversion parameter (see e.g. Bakshi et al. (2003)) which can be estimated given estimates of the risk premium,  $\mu - r$ , and return volatility,  $\sigma$ , obtained from observed returns data. Equation 7, which is repeated here,

$$\nu + \frac{1}{2} \approx \frac{\mu - r}{\sigma^2} + \frac{1}{2} \frac{(\mu - r)^2 + \frac{\sigma^4}{12}}{\sigma^3} \alpha^*,$$

shows that ignoring skewness (the last term) leads to upward bias in the estimate of the price of risk and, hence, of risk aversion. Note that the role of skewness was first highlighted by Kraus and Litzenberger (1976). Intuitively, when agents are risk-averse, and the risk premium is positive, a more negative value of skewness corresponds to an increase in the quantity of risk: the probability of lower returns increases. Then accounting for skewness reduces the price of risk required to fit the observed equity premium and, ultimately, leads to lower estimates of risk aversion in the economy.

The impact of skewness is economically significant. Since 1980, the sample mean and volatility

---

<sup>21</sup>Note that our results contrast with Bakshi and Kapadia (2003). They consider regressions of BSM Delta-Hedged gains on skewness (and kurtosis) and find that it plays a small roll relative to volatility. These regression pooled Delta-Hedged gains across moneyness and are limited to near-the-money options. Looking at Figure 5d, we expect that they produce low coefficients.

of one-year returns is 14.72% and 6.13%, respectively, and the first term of Equation 7 is equal to 20.5. In other words, if risk is summarized by the volatility of market returns, then the equity premium appears too large and leads to excessively high estimates of the coefficient of risk aversion. However, the coefficient of skewness,  $\alpha$ , in the last term is 12.88. For a value of skewness, say, of -1, the estimate of the price of risk is 7.63, less than half than if we ignore the impact of skewness. Moreover, the estimates of skewness we obtain below are often lower than -1. (See also Harvey and Siddique (2000) for evidence that skewness is priced.)

The results presented here show that this important linkage between skewness and the equity premium also hold when using option-based measures of skewness and, furthermore, that the importance of skewness extends to the volatility spread. Note that we provide evidence from two different sources. We first show that conditioning on skewness improves the predictability of future *stock* returns when using the volatility spread. Periods where returns are more skewed also have a higher equity premium and higher volatility spreads. Next we show that conditioning the  $\Delta_t$  of an option on the current estimate of skewness improves delta-hedged gains on out-of-the-money *options*. These options offer more compensation for risk in periods where returns are more skewed. Overall, the evidence suggests that an understanding of the volatility spread and of the equity premium demands an understanding of the determinants of skewness. Moreover, it shows that properly conditioning on implied skewness is key to deciphering the information content of options prices for future returns. In fact, reversing the relationship, and interpreting the volatility spread as the returns on a specific portfolio of options,

$$\begin{aligned} \sqrt{\text{Var}_t^Q [\ln (S_{t+i}/S_t)]} - \sqrt{\text{Var}_t^P [\ln (S_{t+i}/S_t)]} = \\ - \frac{1}{2} \text{Skew}_t^Q [\ln (S_{t+i}/S_t)] \left( E_t^P [\ln (S_{t+i}/S_t)] - E_t^Q [\ln (S_{t+i}/S_t)] \right) \end{aligned}$$

we see that a version of the CAPM conditional on skewness “explains” the returns on the volatility spread portfolio. This offers an answer to the question posed in Carr and Wu (2009) which asks what factor may explain the volatility spread.

Our results complement existing results (e.g. Bakshi and Kapadia (2003), Bollerslev et al. (2008)) where the volatility spread is linked to variance risk being priced. In our model, the asymmetry in returns shifts the risk premium *and* the risk-neutral volatility. This induces the link between the volatility spread and the equity premium. In contrast, Polimenis (2006) and Bakshi and Madan (2006) link the volatility spread to higher order moments of the *historical* distribution.<sup>22</sup> Our results support the relevance of this channel and suggest that an understanding of the volatility spread, and its relationship with compensation for risk, demands an understanding of skewness. In particular, this new stylized fact should help discriminate across competing theories of the observed volatility spread. Our results also complement the results of Bakshi and Kapadia (2003). They test whether

---

<sup>22</sup>Bakshi and Madan (2006) find that measure of *historical* skewness plays a relatively small role in the determination of volatility spread. They did not consider measures of skewness based on options prices.



Delta-Hedged gains for near-the-money options are negative as a test of compensation for volatility risk in option prices. On the other hand, we show the asymmetric impact of skewness on Delta-Hedged gains of options that are away from the money are attributable to the skewness of returns. In particular, this asymmetry cannot be rationalized in a Stochastic volatility model with leverage but with no jumps (see Bakshi and Kapadia (2003)).

## VIII Conclusion

We consider a simple extension of the BSM option pricing model. The Homoscedastic Gamma model allows for arbitrary skewness in the distribution of returns and delivers closed-form option pricing formula at any maturity. We provide a natural change of measure under which returns are HG under the historical and the risk-neutral probability measures.

We first introduce the implied volatility and skewness surface, which we study empirically. This is a new tool that provide a transparent interpretation of variations in the shape and level of the IV curve in terms of skewness. Next, we show that practitioner's variants of the HG models improves upon practitioner's version of the BSM model. This can be interpreted as evidence that expansions around the Gaussian density are not sufficiently flexible to capture the skewness implicit in option prices. More importantly, our model makes explicit the relationship between skewness, the volatility spread, the equity premium and Delta-Hedged gains. The evidence we present support the key role of skewness for option prices and risk premium.

At a first level, this paper provides simple but powerful tools allowing practitioners to monitor and assess the impact of skewness variations on option prices and risk sensitivities. Moreover, the evidence presented here implies that the properties of stock returns volatility and skewness must be considered jointly. While many models have similar implications for the volatility spread, they typically differ in their implications for the skewness implicit in option prices.

## References

- Alexander, C. and L. M. Nogueira (2005). Model-free hedge ratios and scale-invariant models. *Journal of Banking and Finance* 31, 1839–1861.
- Andersen, T., T. Bollerslev, and F. Diebold (2005). Parametric and nonparametric volatility measurement. forthcoming in *The Handbook of Financial Econometrics*.
- Bakshi, G., C. Cao, and Z. Chen (1997). Empirical performance of alternative option pricing models. *The Journal of Finance* 51, 549–584.
- Bakshi, G. and N. Kapadia (2003). Delta-hedged gains and the negative market volatility risk premium. *Review of Financial Studies* 16, 527–566.
- Bakshi, G., N. Kapadia, and D. Madan (2003). Stock return characteristics, skew laws, and the differential pricing of individual equity options. *Review of Financial Studies* 16, 101–143.
- Bakshi, G. and D. Madan (2000). Spanning and derivative-security valuation. *Journal of Financial Economics* 58, 205–238.
- Bakshi, G. and D. Madan (2006). A theory of volatility spreads. *Management Science* 52, 1945–1956.
- Bates, D. (2000). Post-'87 crash fears in the SP500 futures option market. *Journal of Econometrics* 94, 181–238.
- Bates, D. (2005). Hedging the smirk. *Financial Research Letters* 2, 195–200.
- Bates, D. S. (1995). Testing option pricing models. Working Paper 5129, National Bureau of Economic Research.
- Bollerslev, T., G. Tauchen, and H. Zhou (2008). Expected stock returns and variance risk premia. *ERID Working Paper Series* 5.
- Carr, P. and L. Wu (2009). Variance risk premiums. *Review of Financial Studies* 22.
- Christoffersen, P., R. Elkamhi, B. Feunou, and K. Jacobs (2009). Option valuation with conditional heteroskedasticity and non-normality. *Review of Financial Studies*. forthcoming.
- Christoffersen, P., S. Heston, and K. Jacobs (2006). Option valuation with conditional skewness. *Journal of Econometrics* 131, 253–284.
- Christoffersen, P., K. Jacobs, and G. Vainberg (2008). Forward-looking betas.
- Corrado, C. and T. Su (1996). Skewness and kurtosis in SP500 index returns implied by options prices. *The Journal of Financial Research* 19.
- Dennis, P. and S. Mayhew (2000). Risk-neutral skewness: Evidence from stock options. *The Journal of Financial and Quantitative Analysis* 37, 471–493.
- Dumas, B., J. Fleming, and R. Whaley (1998). Implied volatility functions: Empirical tests. *The Journal of Finance* 53, 2059–2106.
- Galai, D. (1983). The components of the returns from hedging options against stocks. *The Journal of Business* 56, 45–54.
- Harvey, R. V. and A. Siddique (2000). Conditional skewness in asset pricing tests. 55, 1263–1295.
- Heston, C. (1993). Invisible parameters in option pricing. *The Journal of Finance* 48, 933–947.

- Jarrow, R. and A. Rudd (1982). Approximate option valuation for arbitrary stochastic processes. *Journal of Financial Economics* 10.
- Jondeau, E. and M. Rockinger (2001). Gram-charlier densities. *Journal of Economic Dynamics & Control* 25.
- Kim, T.-H. and H. White (2003). On more robust estimation of skewness and kurtosis: simulation and application to the SP500 index.
- Kraus, A. and R. Litzenberger (1976). Skewness preference and the valuation of risk assets. *JF* 31.
- León, A., J. Mencía, and E. Sentana (2006). Parametric properties of semi-nonparametric distributions, with applications to option valuation. CEMFI.
- Polimenis, V. (2006). Skewness corrections for asset pricing. Working Paper.
- Potters, M., R. Cont, and J.-P. Bouchaud (1998). Financial markets as adaptive systems. *Europhysics Letters* 41.
- Rompolis, L. and E. Tzavalis (2008). The effects of the risk-neutral skewness on implied volatility regressions. Working Paper.
- Rubinstein, M. and J. Jackwerth (1998). Recovering probability distribution from option prices. *The Journal of Finance* 51, 1611–1631.
- Zhang, J. and Y. Xiang (2005). Implied volatility smirk. *Quantitative Finance* 8, 263–284.

Table II: Summary statistics for strike price and maturity categories. SP 500 call options January 1996 - December 2004.

(a) Summary statistics by moneyness

	Moneyness					All
	<0.95	<0.975	<1	<1.025	>1.025	
Number of Contracts	3343	2418	3859	3077	3809	16506
Average Call Price	28.24	31.80	37.22	47.05	78.85	46.05
Average IV	19.43	19.23	19.36	20.13	22.66	20.26

(b) Summary statistics by maturities

	Contract Month						All
	1	2	3	4-6	7-9	10-12	
Number of Contracts	4303	4016	2377	2822	1726	1167	16506
Average Call Price	36.60	39.53	42.91	51.53	61.95	72.74	46.05
Average IV	20.47	20.24	20.37	20.19	20.15	20.24	20.26

(c) Summary statistics by moneyness and maturities. Number of contracts in brackets.

Months	Moneyness				
	<0.95	0.95 to 0.975	0.975 to 1	1 to 1.025	>1.025
1	[96] 21.39	[398] 18.65	[1104] 18.63	[1172] 19.55	[1533] 22.92
2	[354] 19.80	[668] 18.66	[1113] 19.13	[848] 20.08	[1033] 22.75
3	[461] 19.75	[445] 19.24	[647] 19.78	[406] 20.94	[418] 22.61
4-6	[973] 19.27	[481] 19.48	[504] 20.00	[371] 20.88	[493] 22.39
7-9	[805] 19.18	[262] 20.35	[280] 20.33	[167] 21.26	[212] 22.46
10-12	[639] 19.44	[157] 20.72	[194] 20.99	[89] 21.48	[88] 22.30

Table III: In-sample RMSE

RMSE by moneyness and by maturity in percentage of BSM model's RMSE. BSM is the Black-Scholes Model, HG is the Homoscedastic Gamma Model, P-BSM and P-HG are practitioner's versions of these models where volatility is quadratic in moneyness. P-HG1 is a version where the linear term is zero (i.e.  $\gamma_1 = 0$ ), P-HG2 imposes that  $\beta$  is the risk-neutral skewness (see text) and P-HG3 is unrestricted. SP-BSM and SP-HG are smoothed version of these models where the shape of the quadratic IV curve is constant through the sample. SP 500 call options January 1996 - December 2004.

(a) In-sample RMSE by moneyness

Model	Moneyness					
	$S/X < 0.95$	$0.95 < S/X < 0.975$	$0.975 < S/X < 1$	$1 < S/X < 1.025$	$1.025 < S/X$	All
HG	0.584	0.639	0.649	0.681	0.570	0.368
P-BSM	0.487	0.829	0.901	0.665	0.351	0.350
P-HG1	0.437	0.537	0.536	0.629	0.595	0.278
P-HG2	0.449	0.583	0.579	0.532	0.410	0.234
P-HG3	0.298	0.473	0.471	0.453	0.313	0.135
SP-BSM	0.545	0.812	0.919	0.759	0.477	0.418
SP-HG1	0.562	0.632	0.709	0.712	0.655	0.399
SP-HG2	0.488	0.629	0.709	0.642	0.505	0.312
SP-HG3	0.485	0.667	0.744	0.692	0.469	0.322

(b) In-sample RMSE by maturity

Model	Maturity					
	1	2	3	4-6	7-9	All
HG	0.916	0.729	0.585	0.385	0.569	0.368
P-BSM	0.891	0.735	0.614	0.382	0.526	0.350
P-HG1	0.845	0.697	0.544	0.355	0.420	0.278
P-HG2	0.652	0.593	0.529	0.335	0.437	0.234
P-HG3	0.547	0.450	0.405	0.290	0.305	0.135
SP-BSM	1.018	0.798	0.624	0.379	0.592	0.418
SP-HG1	0.892	0.800	0.669	0.493	0.530	0.399
SP-HG2	0.771	0.647	0.612	0.438	0.505	0.312
SP-HG3	0.916	0.678	0.542	0.335	0.525	0.322

Table IV: Out-of-sample RMSE

Weekly out-of-sample RMSE by moneyness and by maturity in percentage of BSM model's RMSE. Parameters obtained for a given week are held constant to price options observed the following week. BSM is the Black-Scholes Model, HG is the Homoscedastic Gamma Model, P-BSM and P-HG are practitioner's versions of these models where volatility is quadratic in moneyness. P-HG1 is a version where the linear term is zero (i.e.  $\gamma_1 = 0$ ), P-HG2 imposes that  $\beta$  is the risk-neutral skewness (see text) and P-HG3 is unrestricted. SP-BSM and SP-HG are smoothed version of these models where the shape of the quadratic IV curve is constant through the sample. SP 500 call options January 1996 - December 2004.

(a) Out-of-sample RMSE by moneyness

Model	Moneyness					All
	$S/X < 0.95$	$0.95 < S/X < 0.975$	$0.975 < S/X < 1$	$1 < S/X < 1.025$	$1.025 < S/X$	
HG	0.795	0.895	0.877	0.840	0.715	0.657
P-BSM	0.718	0.936	0.999	0.914	0.748	0.676
P-HG1	0.736	0.888	0.869	0.833	0.737	0.621
P-HG2	0.730	0.906	0.892	0.829	0.840	0.658
P-HG3	0.656	0.855	0.876	0.832	0.733	0.568
SP-BSM	0.745	0.930	0.999	0.908	0.683	0.671
SP-HG1	0.774	0.880	0.904	0.860	0.763	0.665
SP-HG2	0.724	0.865	0.895	0.852	0.801	0.639
SP-HG3	0.725	0.885	0.927	0.872	0.696	0.625

(b) Out-of-Sample RMSE by maturity

Model	Maturity					All
	1	2	3	4-6	7-9	
HG	1.059	0.894	0.859	0.715	0.757	0.657
P-BSM	1.069	0.942	0.886	0.727	0.744	0.676
P-HG1	1.023	0.902	0.871	0.718	0.695	0.621
P-HG2	1.264	0.914	0.876	0.708	0.695	0.658
P-HG3	0.996	0.871	0.865	0.717	0.628	0.568
SP-BSM	1.068	0.923	0.864	0.708	0.764	0.671
SP-HG1	1.002	0.933	0.896	0.756	0.728	0.665
SP-HG2	1.055	0.894	0.857	0.725	0.724	0.639
SP-HG3	1.040	0.877	0.855	0.702	0.726	0.625

Table V: Hedging Errors I

Weekly hedging errors by moneyness in dollars. BSM is the Black-Scholes Model, HG is the Homoscedastic Gamma Model, P-BSM and P-HG are practitioner's versions of these models where volatility is quadratic in moneyness. P-HG1 is a version where the linear term is zero (i.e.  $\gamma_1 = 0$ ), P-HG2 imposes that  $\beta$  is the risk-neutral skewness (see text) and P-HG3 is unrestricted. SP-BSM and SP-HG are smoothed version of these models where the shape of the quadratic IV curve is constant through the sample. SP 500 call options January 1996 - December 2004.

(a) 1-week Hedging Horizon

Model	Moneyness					
	$S/X < 0.95$	$0.95 < S/X < 0.975$	$0.975 < S/X < 1$	$1 < S/X < 1.025$	$1.025 < S/X$	All
BSM	0.367	0.200	-0.031	-0.207	-0.390	0.016
HG	0.141	-0.204	-0.212	-0.021	0.177	-0.035
P-BSM	-0.188	-0.124	-0.022	0.127	0.021	-0.046
P-HG1	0.072	-0.103	-0.094	0.042	0.127	0.001
P-HG2	0.085	-0.117	-0.136	0.132	0.174	0.014
P-HG3	-0.028	-0.105	-0.048	0.101	0.135	0.001
SP-BSM	-0.123	-0.122	-0.070	0.048	0.039	-0.054
SP-HG1	-0.154	-0.071	0.086	0.260	0.025	0.023
SP-HG2	0.075	-0.077	-0.096	-0.003	0.024	-0.018
SP-HG3	-0.041	-0.146	-0.112	0.004	0.070	-0.053

(b) 2-week Hedging Horizon

Model	Moneyness					
	$S/X < 0.95$	$0.95 < S/X < 0.975$	$0.975 < S/X < 1$	$1 < S/X < 1.025$	$1.025 < S/X$	All
BSM	0.546	0.219	-0.059	-0.429	-0.817	-0.019
HG	0.182	-0.311	-0.239	-0.076	0.013	-0.082
P-BSM	-0.219	-0.122	0.018	0.071	-0.130	-0.082
P-HG1	0.047	-0.129	-0.043	0.060	-0.039	-0.019
P-HG2	0.131	-0.122	-0.070	0.174	0.168	0.046
P-HG3	-0.030	-0.090	0.003	0.131	0.030	0.002
SP-BSM	-0.176	-0.168	-0.059	0.006	-0.161	-0.114
SP-HG1	-0.231	-0.021	0.252	0.363	-0.138	0.037
SP-HG2	0.086	-0.068	0.034	0.027	-0.136	0.002
SP-HG3	-0.082	-0.250	-0.112	-0.019	-0.049	-0.106

Table VI: Hedging Errors II

Weekly hedging errors by moneyness in dollars. BSM is the Black-Scholes Model, HG is the Homoscedastic Gamma Model, P-BSM and P-HG are practitioner's versions of these models where volatility is quadratic in moneyness. P-HG1 is a version where the linear term is zero (i.e.  $\gamma_1 = 0$ ), P-HG2 imposes that  $\beta$  is the risk-neutral skewness (see text) and P-HG3 is unrestricted. SP-BSM and SP-HG are smoothed version of these models where the shape of the quadratic IV curve is constant through the sample. SP 500 call options January 1996 - December 2004.

(a) 3-week Hedging Horizon

Model	Moneyness					
	$S/X < 0.95$	$0.95 < S/X < 0.975$	$0.975 < S/X < 1$	$1 < S/X < 1.025$	$1.025 < S/X$	All
BSM	0.707	0.383	-0.134	-0.652	-1.212	0.015
HG	0.211	-0.341	-0.345	-0.150	-0.071	-0.116
P-BSM	-0.330	-0.189	-0.111	-0.081	-0.184	-0.197
P-HG1	0.067	-0.120	-0.068	0.024	-0.119	-0.031
P-HG2	0.095	-0.110	-0.062	0.185	0.155	0.037
P-HG3	-0.073	-0.108	0.014	0.066	0.013	-0.029
SP-BSM	-0.315	-0.248	-0.176	-0.067	-0.233	-0.223
SP-HG1	-0.307	0.049	0.325	0.379	-0.250	0.018
SP-HG2	-0.002	-0.024	-0.020	-0.035	-0.233	-0.043
SP-BS3	-0.227	-0.331	-0.216	-0.069	-0.134	-0.211

(b) 4-week Hedging Horizon

Model	Moneyness					
	$S/X < 0.95$	$0.95 < S/X < 0.975$	$0.975 < S/X < 1$	$1 < S/X < 1.025$	$1.025 < S/X$	All
BSM	0.988	0.466	-0.327	-0.930	-1.709	0.022
HG	0.391	-0.274	-0.357	-0.381	-0.391	-0.104
P-BSM	-0.291	-0.197	-0.237	-0.277	-0.463	-0.278
P-HG1	0.187	-0.043	-0.057	-0.155	-0.416	-0.029
P-HG2	0.240	-0.060	-0.063	-0.015	0.070	0.058
P-HG3	-0.028	-0.026	0.024	-0.093	-0.219	-0.046
SP-BSM	-0.243	-0.276	-0.303	-0.284	-0.460	-0.294
SP-HG1	-0.201	0.135	0.346	0.265	-0.502	0.017
SP-HG2	0.126	0.087	0.025	-0.227	-0.483	-0.021
SP-HG3	-0.155	-0.335	-0.255	-0.249	-0.352	-0.249



Table VII: Monthly RMSE

Monthly RMSE by moneyness and by maturity in percentage of BSM model's RMSE. BSM is the Black-Scholes Model, HG is the Homoscedastic Gamma Model, P-BSM and P-HG are practitioner's versions of these models where volatility is quadratic in moneyness. P-HG1 is a version where the linear term is zero (i.e.  $\gamma_1 = 0$ ), P-HG2 imposes that  $\beta$  is the risk-neutral skewness (see text) and P-HG3 is unrestricted. SP-BSM and SP-HG are smoothed version of these models where the shape of the quadratic IV curve is constant through the sample. SP 500 call options January 1996 - December 2004.

(a) Monthly In-Sample RMSE

Model	Moneyness					All
	$S/X < 0.95$	$0.95 < S/X < 0.975$	$0.975 < S/X < 1$	$1 < S/X < 1.025$	$1.025 < S/X$	
HGM	0.737	0.894	0.924	0.774	0.606	0.757
P-BSM	0.615	0.897	0.914	0.713	0.463	0.674
P-HG1	0.628	0.832	0.834	0.728	0.662	0.700
P-HG2	0.631	0.834	0.868	0.713	0.544	0.679
P-HG3	0.540	0.777	0.796	0.661	0.465	0.605
SP-BSM	0.647	0.891	0.919	0.750	0.513	0.700
SP-HG1	0.676	0.862	0.867	0.753	0.740	0.747
SP-HG2	0.641	0.835	0.877	0.752	0.603	0.701
SP-HG3	0.571	0.800	0.808	0.686	0.506	0.632

(b) Monthly Out-of-Sample RMSE

Model	Moneyness					All
	$S/X < 0.95$	$0.95 < S/X < 0.975$	$0.975 < S/X < 1$	$1 < S/X < 1.025$	$1.025 < S/X$	
HG	0.908	1.019	0.996	0.902	0.746	0.913
Q-BSM	0.855	0.988	0.991	0.925	0.761	0.891
Q-HG1	0.880	1.028	0.991	0.904	0.776	0.906
Q-HG2	0.867	1.020	0.996	0.902	0.752	0.897
Q-HG3	0.841	1.001	0.990	0.918	0.764	0.886
SQ-BSM	0.852	0.990	0.990	0.933	0.769	0.892
SQ-HG1	0.898	1.024	0.986	0.903	0.822	0.919
SQ-HG2	0.861	0.977	0.968	0.907	0.789	0.889
SQ-HG3	0.833	0.998	0.979	0.913	0.770	0.881

Table VIII: Delta-Hedged Gains

Time-series average of weekly Delta-Hedged gains disaggregated by moneyness and by maturity. Gains are computed from holding a call option where the exposure to variations in the underlying is hedged through short a position of  $\Delta_t$  unit of the index where  $\Delta_t$  is computed using the BSM or the HG model, respectively. SP 500 call options January 1996 - December 2004.

(a)  $\Delta_t$  computed from BSM


Maturity	Moneyness					All
	$S/X < 0.95$	$0.95 < S/X < 0.975$	$0.975 < S/X < 1$	$1 < S/X < 1.025$	$1.025 < S/X$	
1	-11.08	-5.71	-2.72	-1.38	-1.69	-3.41
2	-7.46	-2.54	-1.27	-0.45	-0.79	-2.27
3	-4.08	-1.36	-0.42	-0.48	-0.50	-1.41
4-6	-2.34	-0.44	-0.13	0.11	0.12	-0.84

(b)  $\Delta_t$  computed from HG

Maturity	Moneyness					All
	$S/X < 0.95$	$0.95 < S/X < 0.975$	$0.975 < S/X < 1$	$1 < S/X < 1.025$	$1.025 < S/X$	
1	-7.12	-4.93	-2.81	-1.57	-0.95	-2.25
2	-3.93	-1.65	-0.90	-0.37	-0.28	-1.53
3	-2.08	-1.02	-0.53	-0.12	0.07	-1.14
4-6	-1.17	-0.52	-0.73	-0.44	-0.65	-0.85

Figure 1: Implied Volatility curves across values of skewness in level (Panel (a)) and in percentage deviation relative to the benchmark (i.e. zero skewness) BSM case (Panel (b)) , The grid covers 41 equidistant values of skewness and moneyness is defined as  $\frac{\ln(S/K)(-r\tau)}{\sigma\sqrt{\tau}}$  to correct for maturity differences.

(a) Level



(b) Percentage Deviation from BSM




Figure 2: Deviations of implied volatility and skewness surfaces from the BSM IV values for different maturity categories. Moneyness is defined as  $\ln(S/K)(-r\tau)$  and maturity groups are defined using settlement dates.

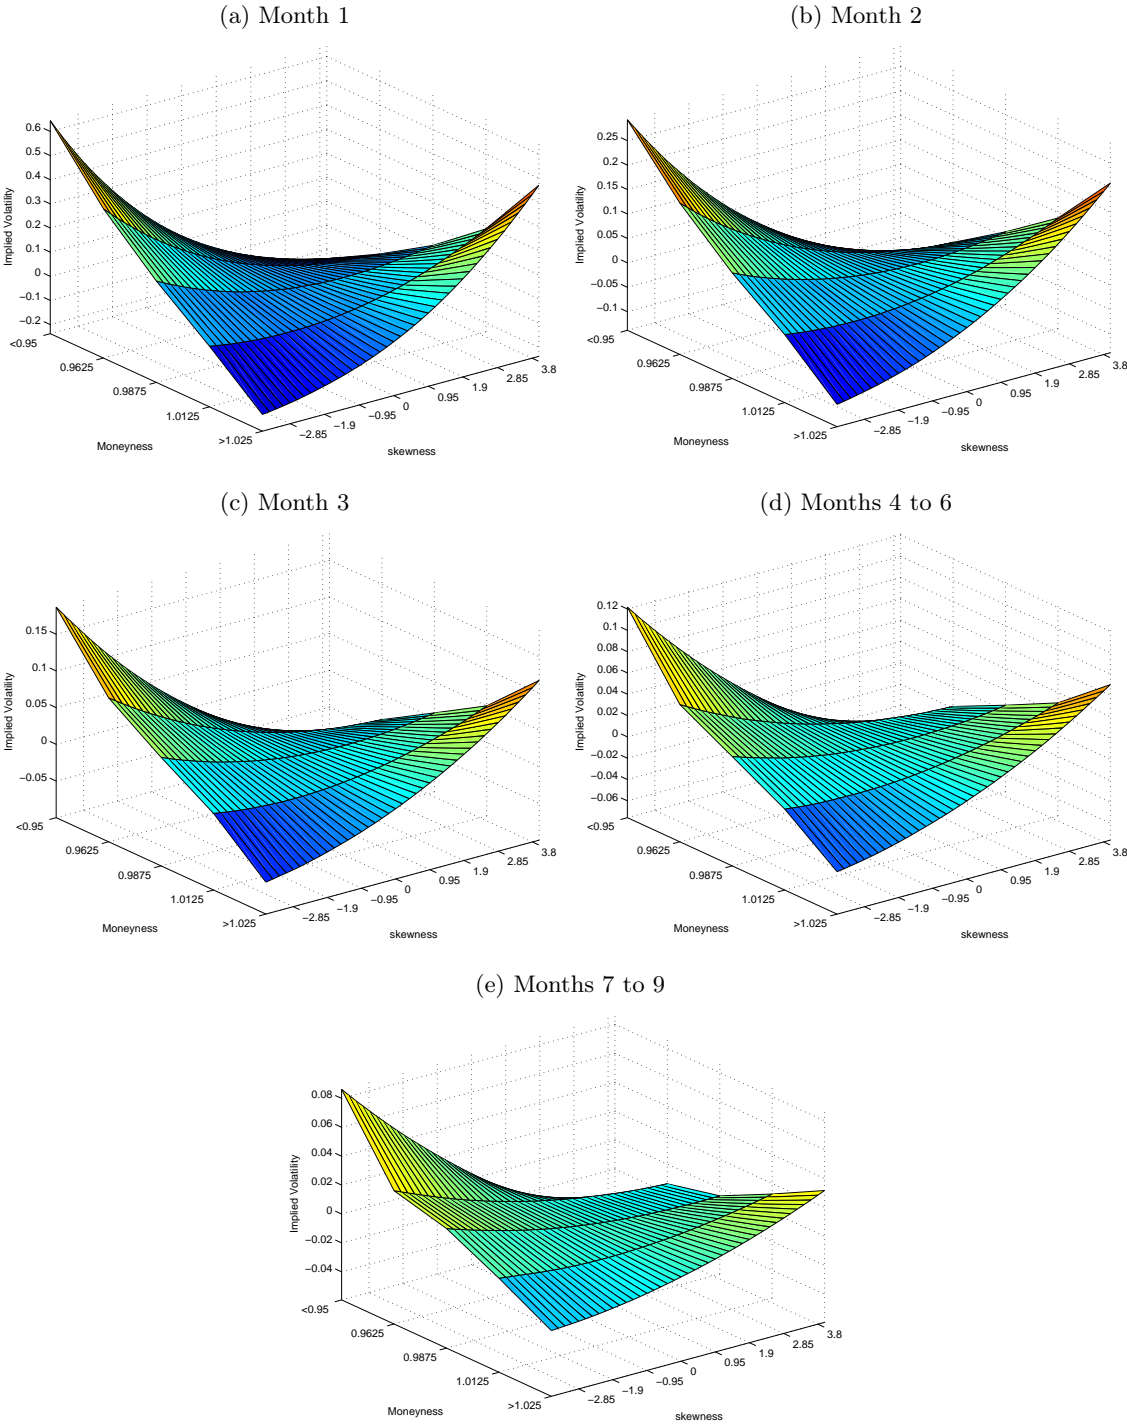


Figure 3: Implied skewness curve for different values of volatility, in percentage deviation from BSM IV values, for different maturity groups. Moneyness is defined as  $\ln(S/K)(-r\tau)$  and maturity groups are defined using settlement dates.

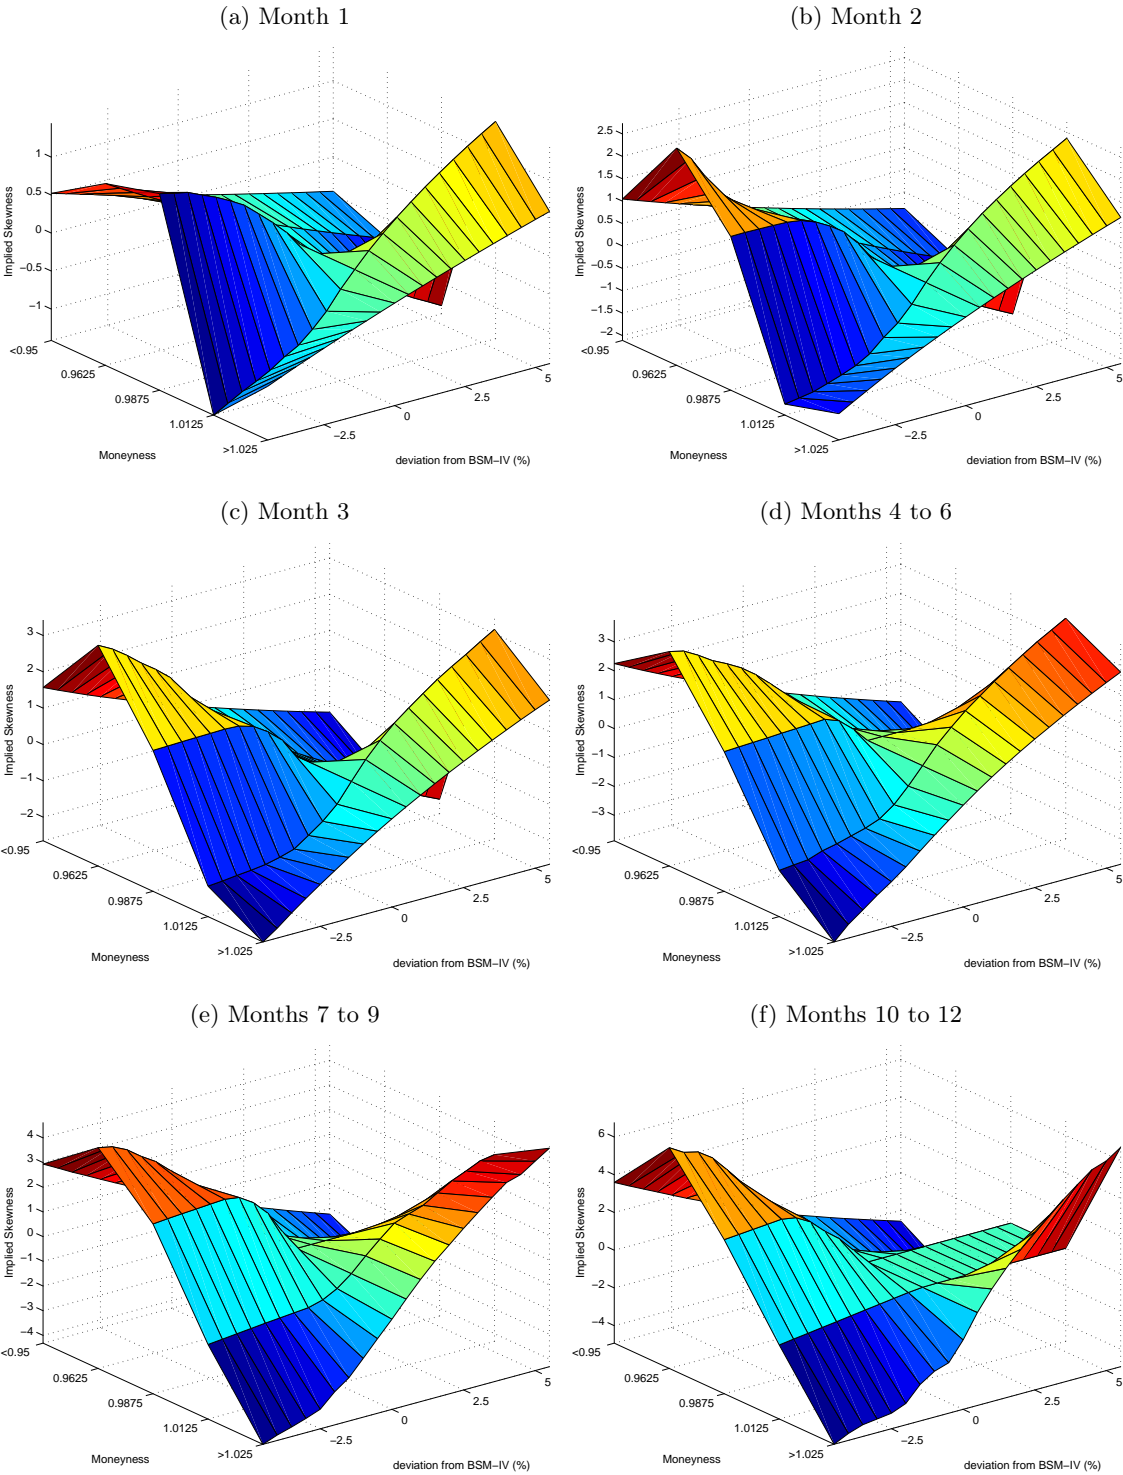




Figure 4: Time-series average of estimates of  $\Theta_t = (\sigma_t, \gamma_{1,t}, \gamma_{2,t})$  from the P-HG3 (unrestricted) model but for different values of skewness. The parameters govern the IV curve:  $\sigma_{i,t} = \sigma_{I0,t}(1 + \gamma_{1,t}\xi_i + \gamma_{2,t}\xi_{i,t}^2)$ .

(a)  $\sigma_{I0,t}$  and Skewness



(b)  $\gamma_{1,t}$  and Skewness



(c)  $\gamma_{2,t}$  and Skewness


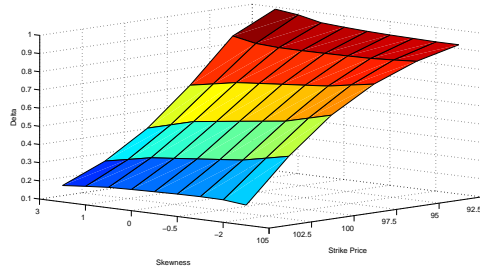



Figure 5: Option prices sensitivities.


(a) First derivative with respect to stock price, in level.



(b) Second derivative with respect to stock price, in level.

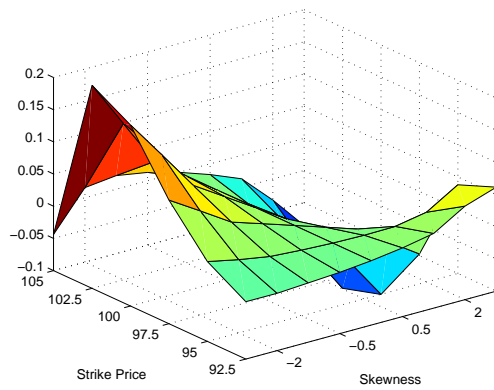


(c) Derivative with respect to volatility, in level.




38

(d) First derivative with respect to stock price, in percentage deviation.



(e) Second derivative with respect to stock price, in percentage deviation.



(f) Derivative with respect to volatility, in percentage deviation.

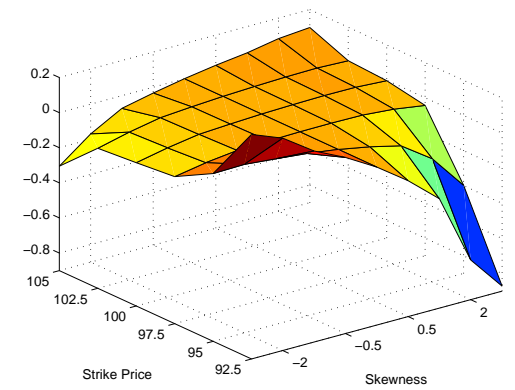


Table IX: Predictability of Excess Returns by Implied Skewness.

The table reports the results of  $n$ -period regressions of returns on the SP500 index in excess of a yield of maturity of  $n$  months:

$$\frac{1}{n} \sum_{j=1}^n \left( r_{M,t+j} - y_{f,t+j}^{(n)} + \frac{IV_t}{2} \right) = a_n + b_n^\top PRED_t + \varepsilon_{n,t+n}.$$

The regressor PRED is a combination of IV-RV and (IV-RV)/IS, where IV and IS are annualized implied volatility and skewness from all option contracts, and RV is the annualized realized volatility. Reported in square brackets and in brackets are respective robust t-statistics for the null that the coefficient is equal to zero, and for the null that the coefficient is equal to  $-2$ . January 1996 to December 2004.

	1	3	6	12	24	36
Constant	-22.19 [-0.65]	-5.43 [-0.20]	-3.50 [-0.12]	-7.14 [-0.24]	-6.93 [-0.24]	-18.96 [-0.70]
(IV-RV)/IS	-3.28 [-2.66] (-1.04)	-2.24 [-2.52] (-0.27)	-2.04 [-2.69] (-0.05)	-2.13 [-3.85] (-0.23)	-1.58 [-2.38] (0.64)	-1.64 [-2.66] (0.57)
Adj. $R^2$	1.85	3.11	5.59	9.72	8.06	11.28
Constant	0.10 [0.00]	2.86 [0.08]	-8.13 [-0.26]	-10.68 [-0.33]	-0.63 [-0.02]	2.31 [0.07]
IV-RV	7.33 [1.76]	6.38 [1.65]	8.11 [3.01]	8.28 [3.40]	4.37 [1.51]	2.12 [0.75]
Adj. $R^2$	-0.03	1.18	5.83	9.72	3.52	-0.11
Constant	-11.78 [-0.33]	-3.23 [-0.10]	-10.59 [-0.34]	-13.93 [-0.44]	-5.15 [-0.16]	-7.55 [-0.25]
IV-RV	-7.46 [-0.93]	-1.53 [-0.25]	4.83 [1.18]	4.63 [1.24]	-1.15 [-0.27]	-5.29 [-1.71]
(IV-RV)/IS	-4.79 [-1.98]	-2.55 [-1.66]	-1.06 [-0.86]	-1.19 [-1.35]	-1.81 [-1.98]	-2.59 [-3.31]
Adj. $R^2$	1.27	2.21	5.55	10.05	7.05	14.06