



BANK OF CANADA
BANQUE DU CANADA

Working Paper/Document de travail
2009-20

The Equity Premium and the Volatility Spread: The Role of Risk-Neutral Skewness

by Bruno Feunou, Jean-Sébastien Fontaine, and
Roméo Tedongap

Bank of Canada Working Paper 2009-20

June 2009

The Equity Premium and the Volatility Spread: The Role of Risk-Neutral Skewness

by

Bruno Feunou,¹ Jean-Sébastien Fontaine,² and Roméo Tedongap³

¹Duke University

²Financial Markets Department
Bank of Canada
Ottawa, Ontario, Canada K1A 0G9
jsfontaine@bankofcanada.ca

³Stockholm School of Economics

Bank of Canada working papers are theoretical or empirical works-in-progress on subjects in economics and finance. The views expressed in this paper are those of the authors. No responsibility for them should be attributed to the Bank of Canada.

Acknowledgements

A previous version of this paper was titled “The Implied Volatility and Skewness Surface”. We thank Peter Christoffersen, Redouane Elkamhi, René Garcia, Scott Hendry, Teodora Paligorova and Jun Yang for comments and discussions. We also thank seminar participants at the Bank of Canada and the Financial Econometrics Lunch Group (Duke University).

Abstract

We introduce the Homoscedastic Gamma [HG] model where the distribution of returns is characterized by its mean, variance and an independent skewness parameter under both measures. The model predicts that the spread between historical and risk-neutral volatilities is a function of the risk premium and of skewness. In fact, the equity premium is twice the ratio of the volatility spread to skewness. We measure skewness from option prices and test these predictions. We find that conditioning on skewness increases the predictive power of the volatility spread and that coefficient estimates accord with theory. In short, the data do not reject the model's implications for the equity premium. We also check the model's implications for option pricing and show that the information content of skewness leads to improved in-sample and out-of-sample pricing performances as well as improved hedging performances. Our results imply that expanding around the Gaussian density is restrictive and does not offer sufficient flexibility to match the skewness and kurtosis implicit in option data. Finally, we document the term structure of option-implied volatility, skewness and kurtosis and find that time-dependence in returns has a greater impact on skewness.

JEL classification: G12, G13

Bank classification: Financial markets

Résumé

Les auteurs présentent le modèle « homoscédastique gamma », ou modèle HG, où la distribution des rendements est caractérisée par sa moyenne, sa variance et un paramètre d'asymétrie. Dans le modèle HG, l'écart entre les volatilités observée et neutre à l'égard du risque est fonction de la prime de risque et du degré d'asymétrie : la prime de risque appliquée aux actions est en fait le double du ratio de l'écart de volatilité à la mesure de l'asymétrie. Les auteurs mesurent l'asymétrie à partir de prix d'options et testent la validité des prédictions de leur modèle. Ils constatent que la prise en compte de l'asymétrie a pour effet d'accroître le pouvoir prédictif de l'écart de volatilité et que les estimations des coefficients sont conformes à ce que prévoit la théorie. En bref, les données ne permettent pas de rejeter les conclusions du modèle concernant la prime relative aux actions. Les auteurs étudient aussi les implications du modèle du point de vue de l'évaluation des options. Ils montrent que le contenu informatif de la mesure de l'asymétrie permet d'améliorer la capacité de prévision du modèle, tant sur échantillon que hors échantillon, ainsi que sa performance en matière de couverture. Les résultats indiquent que le recours à une expansion au voisinage de la densité gaussienne est contraignant et n'offre pas suffisamment de souplesse pour reproduire les degrés d'asymétrie et d'aplatissement qui ressortent du prix des options. Enfin, au terme d'une analyse de la structure d'échéance de la volatilité implicite, de l'asymétrie et de

l'aplatissement, les auteurs concluent que la dépendance temporelle des rendements influe davantage sur l'asymétrie.

Classification JEL : G12, G13

Classification de la Banque : Marchés financiers

I Introduction

We propose the Homoscedastic Gamma model [HG] in which the 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. We follow Christoffersen et al. (2007) and provide a Stochastic Discount Factor [SDF] under which stock returns are HG under both the historical and risk-neutral probability measures. This model delivers a sharp prediction about the relationship between the risk premium, volatility and skewness : the equity premium is equal to twice the *ratio* of the volatility spread to skewness.

The HG model preserves the BSM model parsimony and closed-form option prices. Thus, we measure skewness from option prices. Using these estimates, we perform regressions of SP500 excess returns on the ratio of the volatility spread to skewness. We find that coefficients have the correct sign and magnitude, and that conditioning on skewness improves the predictive power of the volatility spread. In short, the data does not reject the model's restrictions on the equity premium. Reversing the relationship, and interpreting the volatility spread as the returns on a specific portfolio of options, we show that a version of the CAPM conditional on skewness "explains" the returns on the the volatility spread portfolio. This offers a solution to the question posed in Carr and Wu (2008) regarding which factor may explain the variance premium.

The volatility spread has been linked to variance risk (Bakshi and Kapadia (2003), Bollerslev et al. (2008), Carr and Wu (2008)) or to a left-skewed and fat-tailed historical distribution (Bakshi and Madan (2006), Polimenis (2006)).¹ An important implication of this new stylized fact is 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 functions of risk-neutral skewness. In particular, this should help discriminate across competing theories of the observed volatility spread. While different channels have been proposed to explain the volatility spread, they do not have the same predictions for risk-neutral skewness. Clearly, understanding the source of risk-neutral skewness is a key research objective.

As a further check of the HG model, we test its pricing implications for option contracts written on the SP500 index. We consider the simple HG model and variants analogous to the practitioner's version of the BSM model [P-BSM and P-HG]. We interpret these

¹Bakshi 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.

explicitly as expansions around the Gaussian and the HG distributions, respectively. Overall, HG-based models significantly improve in-sample, out-of-sample performance relative to Gaussian-based models with no increase in the number of parameters. They also increase hedging performance at horizons up to 4 weeks.

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. Another way to view these 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, option partial derivatives with respect to the underlying can be computed, model-free, by taking partial derivatives of option prices with respect to strike prices. In practice, however, a parametric model is fitted to observed prices from which derivatives can be imputed. The relative hedging performances of the P-BSM and of the P-HG model imply that the latter offer a better fit of the option price curve across the strike continuum, and a better fit of the true underlying option sensitivities. Still, the improvements come with no increase in implementation costs.

Next, we introduce the implied volatility and skewness *surface*, an extension of the implied volatility curve. Beyond its simplicity and ease of computation, the BSM's implied volatility [IV] curves deliver a transparent comparison of options through time and across strike prices. For traders, this curve is relatively insensitive to variations in the intrinsic value of the options and, thus, in the level of the (volatile) underlying. Rather, it is a measure of an option's time value and provides a direct indication of relative values across strike prices or through time. For researchers, implied volatility curves are viewed as key empirical facts to be matched by options models. Repeating the inversion of the IV curve across values of skewness delivers the implied volatility and skewness surface. The volatility-skewness relationship appears smooth in practice: negative (positive) skewness increases (decreases) the implied volatility of out-of-the-money [OTM] calls and decreases (increases) the implied volatility of in-the-money [ITM] calls. We draw two important conclusions. First, the HG model can restore the symmetry of the observed IV curve. Second, the level of the IV curve also depends on skewness.

Finally, we study the term structure of implied volatility, skewness and excess kurtosis. The HG model requires less data than a non-parametric approach and delivers estimates of risk-neutral volatility and risk-neutral skewness at longer horizons than a non-parametric approach. The evidence suggests that skewness decays at a rate slower than $1/\sqrt{T}$ while kurtosis decays at a rate faster than $1/T$. In other words, the time-dependence structure of returns has a larger impact on risk-neutral skewness than on risk-neutral kurtosis. To our

knowledge, this differential impact of time-dependence on skewness and kurtosis has never been documented.

Related Literature

The stylized observations that IV curves typically display a smile, a skewed smile or a smirk has 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 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 is simplistic 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, approximation of the underlying risk-neutral density often turns 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.² Finally, models based on Gram-Charlier do not provide a change of measure linking the historical and risk-neutral measure.³

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 its link with the index skewness. Also, 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

²Jondeau 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.

³Note 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.

data for future stock betas. However, the pricing or hedging implication of skewness for option prices cannot be handled within this model-free framework.⁴

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 function. Section III presents the data. Section IV perform the regression-based test of the model’s implications for the equity premium and the volatility spread, and discusses the results in the context of equilibrium model. We introduce a practitioner’s analog in Section VI and compare in-sample, out-of-sample and hedging performances of HG and BSM-based models in Section VII. Section V explores the empirical properties of the implied volatility and skewness surface while Section VIII provides estimates of the term structure of volatility, skewness and kurtosis. Section IX concludes.

II The Homoscedastic Gamma Model

This section introduces the Homoscedastic Gamma model for stock returns. The model possesses three crucial properties that makes it a natural choice for our purposes. First, skewness is parameterized directly and is independent of the mean and variance. Second, its density and characteristic function are known in closed-form. Finally, the distribution of returns remains HG for all investment horizons. Also, we show that the return process is HG under both the historical and the risk-neutral probability measures whenever the SDF is exponential in aggregate wealth. This delivers an explicit mapping between moments under each measures. Finally, we obtain closed-form price 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 the gross returns, R_t , over an interval of time Δ , say, follows

$$\begin{aligned} R_{t+\Delta} &\equiv \ln(S_{t+\Delta}/S_t) = \mu^* \Delta + \sqrt{\sigma^{*2}\Delta} \varepsilon_{t+\Delta}^* \\ \varepsilon_{t+\Delta}^* &\sim SG(\alpha^*(\Delta)), \end{aligned} \tag{1}$$

⁴Note, 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.

under the risk-neutral measure where μ^* and σ^{*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 α^* . The SG distribution is defined in terms of the Gamma distribution, $\Gamma(k, \theta)$, as

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

where the scale parameter is fixed to $\theta = 1$. Given that the Gamma definition 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 Δ to reflect the choice of the interval's length. A key simplifying assumption is that the conditional distribution of returns is not-varying. Still, the model could be thought as holding conditionally, with parameters μ_t , σ_t and α_t indexed by time.

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 , that is the return from holding that stock over the period from t until $t + H$. Define $M \equiv \frac{H}{\Delta}$ as the number of time steps over this horizon. It is then easy to see that

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

where the return innovation, $\varepsilon_{t,M}^*$, is given by⁵

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

B Returns Under The Historical Measure

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. In the HG case, the risk-neutral volatility is greater than

⁵This follow directly from the fact that the Gamma distribution is summable.

the historical volatility when the equity premium is positive and skewness is negative. Also, the volatility spread increases with the equity premium and whenever returns become more left-skewed. 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

$$R_{t+\Delta} \equiv \ln(S_{t+\Delta}/S_t) = \mu \Delta + \sqrt{\sigma^2 \Delta} \varepsilon_{t+\Delta}, \quad (3)$$

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

$$M_t = \exp(-\nu(\Delta) \varepsilon_t + \Psi(\nu(\Delta))), \quad (4)$$

for some ν and where Ψ 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 ν , as stated in the following proposition.

Proposition 1. *If stock returns follow Equation 3 and if the Stochastic Discount Factor belongs to the class defined by Equation 4 for some ν , then, this SDF defines an Equivalent Martingale Measure for discounted stock prices if and only if*

$$\nu(\Delta) = -\frac{2}{\alpha(\Delta) \sqrt{\sigma^2 \Delta}} + \frac{g(\Delta)}{g(\Delta) - 1}, \quad (5)$$

where

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

See the Appendix for all proofs. This is a direct application of results from Christoffersen et al. (2007). 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 4.

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 3 and if the Stochastic Discount Factor is as in Equation 4 for ν given in Proposition 1 then stock returns are given by Equation 2 and 3 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.⁶

To see the relationship between ν and skewness, consider a first-order expansion of Equation 5 around $\alpha(\Delta) = 0$. For small deviations around the symmetric case we have

$$\nu(\Delta) \approx \frac{\mu - r}{\sigma^2} + \frac{1}{2} + \frac{(\mu - r)^2 + \frac{\sigma^4}{12}}{\sigma^3}\beta(\Delta), \quad (6)$$

Note that $\nu(\Delta)$ tends toward the usual result, $\frac{\mu - r}{\sigma^2}$, when skewness approaches zero. Then, as expected, ν 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.⁷ Taking expectations and re-arranging reveals the following important restriction between the equity premium,

⁶One 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.

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

the volatility spread and the risk-neutral skewness,

$$E_t^P[\ln(S_{t+\Delta}/S_t)] - E_t^Q[\ln(S_{t+\Delta}/S_t)] = -2\frac{\sigma^* - \sigma}{\alpha^*\sqrt{\Delta}}. \quad (7)$$

Therefore, regressions of excess returns on the ratio of the volatility spread to skewness should be more informative than the spread itself. Moreover, the constant is zero and the predicted value for the coefficients is -2. This provides a simple test of the importance of skewness. In the HG model the volatility spread is solely due to the presence of skewness and not to volatility being priced. Since the volatility spread increases when skewness is more negative, the equity premium increases when returns become more left-skewed. We test these implications explicitly below.

C Option Prices

We are now ready to provide a closed-form solution for European style contingent claims on a stock. 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. That is,

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

As usual, the solution is a 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 2, then the price of a European call option is*

$$C_t(K, H) = S_t C_{1,t} - e^{(-rH)} K C_{2,t}, \quad (8)$$

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

$$C_{1,t} = P\left(\frac{H}{\beta(\Delta)^2}, d_1(\Delta)\right) \quad (9)$$

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

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

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

$$C_{2,t} = Q\left(\frac{H}{\beta(\Delta)^2}, d_2(\Delta)\right), \quad (12)$$

The functions $P(a, z)$ and $Q(a, z)$ are the regularized gamma functions⁸ 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 functions⁹ 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 2 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.

Next, we introduce a second sample that group option prices at the monthly frequency. This reduces the noise in the estimates of volatility and skewness we use in excess returns regressions. Another benefit of this approach is that it ensures enough observations to estimate our model in each maturity group. This allows us to draw the implied volatility

⁸We 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 .

⁹Note that we have $P(a, z) + Q(a, z) = 1$, which is a convenient property when computing derivatives (see below).

and skewness surface in different maturity groups and, as well, to obtain a term structure of skewness and volatility. To group observations, we use settlement dates rather than calendar months. Since each contract settles on the third Friday of a month, we group all observations intervening between two successive settlement dates.¹⁰ All weekly observations occurring within such a sub-period can be unambiguously attributed to one maturity group.¹¹ Note that settlement dates follow a regular pattern through time: contracts are available for 3 successive months and then for the next 3 months in the March, June, September, December cycle. This leads to maturity groups with 1, 2 or 3 months remaining to settlement and then between 3 and 6, between 6 and 9, and between 9 and 12 months remaining to settlement,¹²

Table I 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 level of the 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 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 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 options are deep in-the-money, since the exchange does not regularly issue short horizon out-of-the-money 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.¹³

¹⁰These subperiods have varying length depending on the (calendar) months they cover.

¹¹Take any contract, on any observation date. This contract is assigned to the 1-month maturity group if its settlement date occurs on the following third-Friday, to the 2-month group if it occurs on the next to following third-Friday, etc.

¹²Within a given month, and within a given maturity group, the same contract (i.e. same strike price) is observed with successively shorter maturities. However it is priced consistently under the null of i.i.d. returns innovations throughout the month.

¹³This highlights the importance of using a model which can handle maturity differences. In particular, model based on density approximation are not robust to this composition effect.

Finally, Panel (a) of Figure 1 presents the number of available observations for each day, which averages around 40 and typically ranges between 20 and 50 contracts. Panel (b) decomposes this number and presents the proportion of contracts in each moneyness category. An ongoing extension of this paper is to include both call and put options and obtain a more balanced sample across moneyness as well as a better coverage of longer maturities.

IV The Volatility Spread And The Equity Premium

A Model's Implications

The relationship between skewness and volatility runs deeper than a simple redistribution of probability mass across the support of returns. That is because the impact of the representative investor's preference on the risk-neutral volatility depends on skewness. In the particular case where the representative SDF can be approximated by the exponential-shift given in Equation 4 we have a tight link between the price of risk, the volatility spread and skewness. After some manipulation of Equation 7, we obtain

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

for an investment horizon i and where $r^{(i)}$ is the risk-free rate for that horizon and ω_t is the Jensen adjustment term.¹⁴ 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 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 coefficients should be -2.

B Aggregating Data

We obtain estimates of risk-neutral volatility and skewness from option data. Estimates of skewness are noisy in weekly data. This is in part due to the number of option prices available each week and, also, to the sampling frequency. One simple solution is to group price observations at the monthly level where we define a month as the period between successive expiration dates which occur every third Friday (See Section III). Within each month, the data consists of repeated observations of the same contracts over a period of 4 (or 5) weeks.¹⁵ This implicitly assumes i.i.d. return innovations throughout a month, which

¹⁴This 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, $\frac{1}{2}\sigma^2$.

¹⁵Some contracts are not observed each Wednesday within a month. New contracts become available to participants as the index moves away from the range of available strike prices. Some contracts are not

is consistent with the model and reasonable over this short time span. It also implies that the maturity *date* of each contract is constant throughout each month and, thus, that the skewness estimate pertains to a set of contracts that mature at fixed maturities. Finally, we measure the historical volatility using the observed realized volatility.

We estimate our preferred version of the model each month through minimization of squared pricing errors.¹⁶ Figure 2 presents the time series of our volatility estimates (Panel (a)) and of our skewness estimates (Panel (b)). Skewness typically varies around -1 but dipped close to -2.5 in the summer of 1998 and in the second half of 1999, and slightly below 1.5 in the Fall of 1996 and the Spring of 2004.

C Implied Skewness And The Risk Premia

Table II 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.¹⁷ 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 volatility to skewness ratio. The results from the regressions are presented in Table II. Since volatility and the ratio of the volatility spread to skewness are highly correlated, the coefficients become unreliable, even changing signs in many cases for the volatility spread. However, their combined predictive power does not rise above that of the volatility to skewness ratio, further supporting the implications of the model.

available each week because they were excluded from the weekly sample due to liquidity concerns.

¹⁶Specifically, we estimate a restricted version of the practitioner's HG model that allow for kurtosis but maintain the identification of the risk-neutral volatility and skewness (See Section VII). As a robustness check (not reported) we repeated the exercise using skewness estimated from the simple HG model presented above. The results are not qualitatively different.

¹⁷Precisely, 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

D Discussion

We can also 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, ν 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, σ , obtained from observed returns data. Equation 6, 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. 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, unbiased, estimates of risk aversion in the economy.

Note that the effect of skewness is economically significant. Since 1980, the sample mean and volatility of one-year returns is 14.72% and 6.13%, respectively, and the first term of Equation 6 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, α , 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.

The results shows the linkages implied by the HG model between the equity premium, the volatility spread and the skewness hold (Equation 7). This suggests that an understanding of the volatility spread and of the equity premium demands an understanding of the determinants of risk-neutral skewness. Moreover, it shows that properly conditioning on implied skewness is key to deciphering the information content of options prices for future returns. Recall that in this model the volatility spread is due to the presence of skewness. In fact, reversing the relationship, and interpreting the volatility spread as the returns on a specific portfolio of options,

$$\sqrt{h_t^*} - \sqrt{h_t} = \frac{1}{2\alpha_t^*} \left(E_t^P [\ln (S_{t+i}/S_t)] - E_t^Q [\ln (S_{t+i}/S_t)] \right)$$

we see that a version of the CAPM conditional on skewness “explains” the returns on the

the volatility spread portfolio. This offers a solution to the question posed in Carr and Wu (2008) which asks what factor may explain the volatility spread.

Our results contrast with existing results (e.g. Bakshi and Kapadia (2003), Bollerslev et al. (2008)) where the spread is linked to variance risk being priced. In our model, the asymmetry in returns shifts the price of risk *and* the risk-neutral volatility. This induces the link between the volatility spread and the equity premium. Similarly, Polimenis (2006) and Bakshi and Madan (2006) link the volatility spread to higher order moments of the *historical* distribution. In particular, Bakshi and Madan (2006) conclude that the historical skewness plays a relatively small role in the determination of volatility spread. They did not consider risk-neutral skewness. One important implication is that understanding the source of risk-neutral skewness is a key research objective. From the tight linkages we uncover, we conclude that an understanding of the volatility spread, and its relationship with the compensation for risk, demands an understanding of risk-neutral skewness. In particular, this new stylized fact should help discriminate across competing theories of the observed volatility spread.

Finally, one drawback is that mis-specification of the parametric model affects the results through a measurement errors problem. In an ongoing extension of the paper, we check the robustness of the results to the use of non-parametric measures and increase the efficiency of the test through GLS regressions.

V 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. This allows us to easily measure both the volatility and skewness implicit in option prices. This is important because while alternative parametric (e.g. ARCH) and non-parametric (e.g. Realized Volatility) measures of volatility exist¹⁸, the lack of robustness of the usual sample skewness estimator is well known (e.g. Kim and White (2003)).

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

¹⁸See 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.

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. Moreover, the level of implied volatility or, alternatively, the implied volatility of at-the-money options is sensitive to the choice of the skewness parameter. In particular, this implies that empirical studies of the volatility spread are based on 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 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, altogether, 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. Below we analyze the results within each maturity category, and in the aggregate.

B Impact Of Skewness on Implied Volatility Curves

This section traces the average implied volatility and skewness surface across time and maturity groups. This provides a smoother picture of the impact of skewness on the IV

curve. The average volatility-skewness surface is given in Figure 3 in level (Panel (a)) and in percentage deviations from the the benchmark BSM IV curve (Panel (b)). Panel (a) display 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 -1.¹⁹ Clearly, the HG model is sufficiently flexible to capture the skewness implicit in option prices. 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 value 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.

C Results For Different Option Maturities

Next, Figures 4(a)-(e) present implied volatility and skewness surfaces within different maturity groups while Figures 5(a)-(e) report the same results but in percentage deviations from BSM values. Note first that when skewness is zero, which corresponds to the BSM case, we observe sharp differences between implied surfaces at different 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 eventually becomes smirked. For negative values of skewness, and for any maturity, the IV curve is tilted toward a symmetric smile. For contracts maturing at the next settlement date, small negative values of skewness appears sufficient to establish a symmetric smile. As we increase maturity, however, more negative values are necessary. Looking at Figure 5 we see that the impact of a given variation in skewness is much lower for the longest maturity.

D Impact Of Volatility On Implied Skewness

Figures 6(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.

¹⁹The 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.

However, the impact of volatility on implied skewness is asymmetric and highly nonlinear on both side 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, 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.

VI 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 may due to excess kurtosis.²⁰ In this section, we introduce the practitioner’s variants of the BSM model [P-BSM] and of the HG model [P-HG]. These capture deviations from risk-neutral distributions by modeling volatility as a quadratic function of moneyness. That is, in the P-BSM case,

$$\sigma(\xi) = \sigma_{I0}(\alpha, \kappa)(1 + \gamma_1(\alpha, \kappa)\xi + \gamma_2(\alpha, \kappa)\xi^2)$$

where ξ is moneyness and α and κ are the skewness and excess kurtosis of the risk-neutral distribution. Note that $\sigma(0) = \sigma_{I0}$ and that moneyness is defined as

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

to correct for maturity differences where S, K, τ and $\bar{\sigma}$ denote the stock price, the strike price, the time to maturity and the average implied volatility, respectively.

In the following, we first document that estimates of σ_{I0} , γ_1 , and γ_2 vary when we allow for skewness. This contrast with the usual interpretation of γ_1 as a measure of skewness. Next we show that the P-HG model is justified as an expansion around the HG

²⁰In contrast with the Gaussian case, the kurtosis of the HG distribution varies with parameter values. Its kurtosis is proportional to the square of the skewness.

density. This results is the analog of the justification of the P-BSM model as an expansion around the Gaussian density. However, we provide restrictions on parameters of the IV function such that we can recover estimates of α and κ .

A Unconstrained IV Curves

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. σ_{I_0} , γ_1 and γ_2) are usually interpreted in terms of the variance, skewness and kurtosis of the true underlying risk-neutral distribution.²¹ Zhang and Xiang (2005) argue that in the Gaussian case and up to a first-order approximation $\sigma_{I_0}(\beta, \kappa)$ is linear in the risk-neutral volatility, $\gamma_1(\beta, \kappa)$ is linear in skewness, and $\gamma_2(\beta, \kappa)$ is linear in kurtosis. For these reasons, the coefficients of the IV function are usually left unrestricted at estimation.

To evaluate empirically the impact of skewness on estimated IV curves, we fix the value of β and estimate the P-HG model at each date by minimizing squared pricing errors. We repeat the exercise for different values of skewness. We then average unconstrained estimates of $\sigma_{I_0,t}$, $\gamma_{1,t}$ and $\gamma_{2,t}$ through time and trace their relationship with skewness. Figure 7 presents the results. Panel (a) presents average estimates of σ_{I_0} . For contracts maturing at the next settlement date, at-the-money implied volatility is 20% on average when skewness is zero. When (absolute) skewness increases, estimates of at-the-value volatility increases as high as 23%. Intuitively, shifting some probability mass toward one side 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 β : both parameters are measures of the underlying skewness. Finally, Panel 8c shows that $\hat{\gamma}_2$ also varies (non-linearly) substantially with skewness.²² For each parameter, the variations with skewness are stronger for shorter maturities.

The strong dependence between the skewness and all parameters of the IV curve implies that information on the underlying risk-neutral moments is shared across parameters. Furthermore, the fact that estimates of β and of γ_1 are (linearly) correlated implies that they

²¹ σ_{I_0} controls the level of the curve, $\gamma_1(\beta, \kappa)$ controls the asymmetry of the curve and $\gamma_2(\beta, \kappa)$ controls its steepness.

²²This contrasts with the theoretical results of Zhang and Xiang (2005). 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, their linearization strategy may lead to a poor approximation (see below).

are poorly identified. For our purposes, we need to find restrictions on σ_{I0} , γ_1 and γ_2 such that only $\hat{\beta}$ can capture the risk-neutral skewness. Absent these restrictions, parameters of the IV curve capture some of the asymmetry in the underlying distribution leading to biased estimates of β . Note that merely imposing $\gamma_1 = 0$ does not identify an estimator of β with skewness. The following section introduces a framework which will ultimately lead us to the desired restrictions. Under these restrictions, any deviation from a flat IV curve can only be linked to deviations of κ . This sharp identification of parameters in term of the underlying moments is useful to understand the relative importance of skewness and kurtosis in matching observed prices. More importantly, the unambiguous identification of skewness is necessary to provide a measure of the risk premia from implied volatility and skewness.

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 $\hat{\beta}$ 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. As a by-product, we obtain an estimator of the kurtosis in excess of the Gamma distribution.

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 δ^* is a risk-adjustment term, y is a random variable with mean zero, unit variance, skewness, $\frac{2\beta^*}{\sqrt{T}}$ and kurtosis, λ_2^* . If y is normally distributed, then $\delta = \frac{\sigma^{*2}}{2}$. We allow for non-normality beyond the HG and assume that the probability density of y is given by

$$f(y) = g(y) + \frac{\lambda_2^* - \frac{6\beta^{*2}}{T}}{4!} \frac{d^4 g(y)}{dy^4}, \quad (13)$$

where $g(y)$ is the standardized gamma density. This is a one-term Edgeworth expansion around the standardized gamma distribution. This approach captures fat tails in excess of the Gamma distribution but ignores deviations beyond the fourth moment. Our objective

here is to allow for non-trivial implied volatility and skewness surface due to excess kurtosis and derive explicitly the function $\sigma_{I0}(\kappa)$, $\gamma_1(\kappa)$ 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 13, then the price of a call option, $C^*(K, T)$, with maturity T , underlying price S_0 and strike price K is*

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

when $\beta < 0$ and

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

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

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

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 β delivers a convergent estimate of the risk-neutral skewness β^* . 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

$$\begin{aligned} C^*(K, T) &= C(K, T) \\ \frac{\partial C^*(K, T)}{\partial K} &= \frac{\partial C(K, T)}{\partial K} \\ \frac{\partial^2 C^*(K, T)}{\partial K^2} &= \frac{\partial^2 C(K, T)}{\partial K^2} \end{aligned}$$

when evaluated at-the-money (i.e. $K = S_0 e^{rT}$). These restrictions are given in the appendix but note that they are trivially satisfied whenever $\kappa = 0$ since in this case the HG model is true and the implied volatility-surface is flat. Of course this corresponds to the case $\sigma_{I0} = \sigma$, $\beta = \beta^*$ and $\gamma_1 = \gamma_2 = 0$. We linearize the restrictions around this point (i.e. $\kappa = 0$) and obtain

$$\frac{\sigma_{I0} - \sigma}{\sigma} = A_1(\sigma, \beta)\kappa \tag{14}$$

$$\gamma_1 = B_1(\sigma, \beta)\frac{\sigma_{I0} - \sigma}{\sigma} + B_2(\sigma, \beta)\kappa \tag{15}$$

$$\gamma_2 = C_1(\sigma, \beta)\frac{\sigma_{I0} - \sigma}{\sigma} + C_2(\sigma, \beta)\gamma_1 + C_3(\sigma, \beta)\kappa, \tag{16}$$

where the coefficients are given in the appendix.²³ Then, small deviations of the underlying density from a HG distribution leads to deviations from a flat implied volatility and skewness surface. This highlights the impact of excess kurtosis on the estimates of at-the-money implied volatility, of γ_1 and of γ_2 . It also makes clear that deviations from a flat implied volatility and skewness surface are only due to deviations of excess kurtosis from zero. Also, up to a first order, the impact of excess kurtosis on at-the-money implied volatility feeds back on estimates of γ_1 and γ_2 . More importantly, these restrictions ensure that β corresponds to the risk-neutral skewness and that the practitioner's HG model conforms to the true returns density.

VII 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 performance. The improvements are robust if we impose the identification of the skewness parameters, as discussed in the previous section. Indeed, the improvements remain when

²³We 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.

the only deviation from the simple HG model is a constant adjustment to kurtosis through time. In other words, a fixed implied volatility and skewness surface combined with variations in skewness is sufficient. 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.

We evaluate the basic HG model as well as three different versions of the P-HG model: P-HG1, P-HG2, P-HG3. The first version imposes the simple restriction that $\gamma_1 = 0$. The second model, P-HG2, imposes the restrictions derived in the previous section. Finally, P-HG3 is unrestricted. We also introduce a “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 γ_1 is zero but also holds γ_2 constant through time. Next, SP-HG2 still imposes that all deviations from a flat IV curve are linked to the value of κ but the latter is held constant through time. Finally, SP-HG3 imposes the following structure on the IV curve,

$$\sigma_t(\xi) = \sigma_{I0,t}(1 + (\gamma_{10} + \gamma_{11}\beta_t)\xi + (\gamma_{20} + \gamma_{21}\beta_t)\xi^2).$$

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

A In-Sample RMSE

A.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 course, this is simply a reflection returns begin non-Gaussian, as reported elsewhere in the literature.

A.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 no increase in the number of parameters. Moreover, 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 in choosing risk-neutral skewness and kurtosis.

Next, although the naive $\gamma_1 = 0$ restriction may seem reasonable, it fails in practice by yielding larger RMSEs. Comparing models, we see that the improvement of the P-HG2 over the P-HG1 is substantial and is obtained in short maturity, out-of-the-money options. Finally, with one more parameter, the P-H3 offers much lower in-sample RSME (14%) than any other model and yields substantial improvements across all moneyness and maturity categories.

A.3 Smoothed Coefficients

As expected, the various smoothed models yield lower in-sample performance, although note that the SP-HG2 model still improves (31%) over the P-BSM model. These models retain the variability of the skewness parameters from date to date but impose a constant shape on the IV curve through the sample. Overall their performance suggests that time-variation in the shape of the IV curve may be important in-sample. However, out-of-sample results in the next section show that this result is not robust out-of-sample, indicating a relatively minor role for information beyond the third moment.

B *Out-of-sample RMSE*

The good performance of models based on the HG distribution could be due to over-fitting and might not hold out-of-sample. To check this, we perform the following out-of-sample exercise. First, we estimate each model's parameter from options in a given week. We then fix these parameters and evaluate the model's ability to price options observed in the following week. Table IV presents RMSE for each model across strike prices (Panel (a)) and across maturities (Panel (b)).

The relative out-of-sample RMSE increases for all model. This indicates that some of the deviations from the Gaussian case are transitory. The lowest RMSE is now 57%, obtained for the P-HG3 model, while the worst result is 68%, obtained for the P-BSM model. Moreover, keeping the shape of the IV curve constant (i.e. SP-BSM model) offers similar results (67%). This is another piece of evidence that the practitioner's version of the BSM model does not properly fit the (persistent) skewness and kurtosis present in the data. On-the-other hand, for the purpose of measuring risk-neutral skewness, the performance of the P-HG2 (66%) model is promising.

Strikingly, the SP-HG2 model, which fixes excess kurtosis through the sample, actually improves out-of-sample RMSE (64%) over its the more flexible P-HG2. Although the P-HG2 model is more flexible, it induces large out-of-sample errors for short-maturity, out-of-the-money options. Some of the variations in excess kurtosis required to match (in-sample) option prices in this category seem to be transitory, degrading the out-of-sample performance of the model. A similar remark holds for the P-BSM model. Most of the variations in the shape of the IV curve do not translate into improved out-of-sample pricing. The models differ because in the HG framework we know how to restrict unnecessary variation in kurtosis while retaining the flexibility in skewness. Finally, the P-HG3 model remains the best performer out-of-sample.

C Hedging Errors

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. Again, we find that the unrestricted P-HG3 model performs the best. However, the SP-HG2 model, where skewness is separately identified, offers the next best performance, highlighting, again, the value of theoretically sound restrictions.

C.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 appendix. These derivatives depend on the direction of asymmetry and everywhere the symmetric case (i.e. $\beta = 0$) 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 8 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 options. Also, the sensitivity of Delta to skewness

is large for strike prices that are near the money. 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 becomes monotonous in the strike price and as much as 50% lower than when skewness is zero for in-the-money options and 50% higher for out-of-the-money options. Finally, skewness has a similar, 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 and increases by nearly 20% for in-the-money options.

Clearly, ignoring the impact of skewness can lead to large hedging errors. The results below will confirm this. Even though the Practitioners' BSM [P-BSM] can in principle capture some of the risk-neutral skewness implicit in option prices, either its lack of flexibility, or its inability to identify the separate impact of skewness leads to substantially larger hedging errors.

C.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.²⁴ 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, although the BSM model appears to perform well, its overall mean hedging errors averaging 1.6 cents, this hides important disparities across maturities with average hedging errors ranging from 36.7 cents for out-of-the-money options to -0.39 cents for in-the-money options. Moreover, the P-BSM model has higher overall hedging errors (-4.6 cents) with substantial average errors (-18.8 cents) for the lowest strike prices. The poor performance of the BSM-based models does not come as a surprise but it highlights the importance of considering hedging errors across moneyness. When considering the overall mean and the dispersion of hedging errors across maturities, the best performing models are variants of the P-HG model. In particular the identification restrictions for skewness appear to perform well when kurtosis is held constant through the sample. The SP-HG2 model has low hedging errors and mean errors remain below 10 cents across moneyness category. The P-HG1, P-HG2 and P-HG3 models offer

²⁴This abstracts from the hedging errors due to discrete adjustments. See Galai (1983) for details.

marginally lower but similar performances.

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. Most interestingly, 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 P-HG3 model is still performing well. Its overall average at 3-week horizon is now -2.9 cents and mean hedging errors range from -10.8 cents to 6.6 cents while the SP-HG2 still performs similarly well except for the deep-in-the-money category. These two models still stand out at the 4-week horizon.

D 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 the frequent observation that the P-BSM model offers sufficient flexibility and that a higher number of parameters leads to over-fitting. Indeed, option prices based on the HG distribution offer better performance than the P-BSM with the 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.

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 the delta and gamma exactly from non-parametric derivatives. In practice, however, some parametric model is fitted to observed prices from which derivatives can be imputed. The relative hedging performances of the P-BSM and the P-HG model 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 and kurtosis by Gaussian-based expansions translate in inaccurate option sensitivity measures and larger hedging errors relative to approximations based on the Gamma density.

For our purposes, the performance of the SP-HG2 model is promising. Indeed, imposing the additional restriction that excess kurtosis is constant yields the next to best out-of-

sample and hedging performances. In contrast, fixing the shape of the IV curve leads to lower performances for the P-BSM model and simpler restrictions on the individual parameters of the IV curve do not help. It is interesting to note that the estimate of κ is -0.042. Given that kurtosis is the square of skewness in the HG model, the negative sign allows the algorithm to use more asymmetry to fit the data. This adjustment is significant: to keep kurtosis constant but with κ equal to zero, skewness would have to be reduced (close 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.

VIII Term Structure Of Volatility, Skewness and Kurtosis

Section V presented the trade-off between volatility and skewness when fitting option data. One important observation is that a different value of skewness was required to restore the symmetry of the IV curve. This suggests that the risk-neutral distribution converges at slower rate than implied by the Central Limit Theorem. Since returns innovations are i.i.d. the model predicts that skewness of returns decreases with the square root of horizon. If this is verified in the data, estimates of skewness should not vary with the maturity of options. Otherwise, the term structure of implied skewness reflects a degree of dependence implicit in option prices.

An important advantage of our parametric approach is that we can obtain estimates of risk-neutral moments at much longer horizons than is usually the case with non-parametric methods. We estimate the term structure of volatility, skewness and kurtosis using the SP-HG2 model discussed above. We minimize pricing errors then separately for each maturity (1, 2 and 3 months, and then from 4 to 6 and from 7 to 9 months. See Section III). Figure 9 presents the results.

Figure 10a presents the average implied volatility for each maturity. The time-series average rises from close to 21.4% for the next settlement month to 21.8% at a maturity of 3 months. Thereafter, implied volatility remains more or less flat. Figure 10b presents results for (negative) the implied skewness. In contrast with implied volatility, implied asymmetry rises sharply for all maturities we consider. Figure 10c shows the term structure of (minus) the implied excess kurtosis. Perhaps surprisingly, excess kurtosis relative to the HG distribution decreases with maturity. Overall, the term structure evidence indicates a substantial degree of dependence between the distributions of expected returns. However, the impact of dependence appears to have a much greater impact on implied skewness than

on other moments. While the term structure of implied volatility appears to flatten out beyond a maturity of 3 months, implied excess kurtosis actually decreases with maturity. Recall that kurtosis is proportional to the square of skewness in the HG distribution. Then, the pattern of decreasing excess kurtosis suggests that the dependence structure of returns has a lesser impact on kurtosis than on skewness. The evidence suggests that skewness decays at a rate slower than $1/\sqrt{T}$ while kurtosis decays at a rate faster than $1/T$. To our knowledge, this differential impact of time-dependence on skewness and kurtosis has never been documented.

IX Conclusion

We provide a simple extension of the BSM option pricing model. The Homoscedastic Gamma model, which 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. An important implication is that the relationship between the equity premium and the volatility spreads is conditional on current expected risk-neutral skewness. It is the ratio of the volatility spread to skewness that predicts excess returns. Empirically, we find coefficients that correspond to implications from the model. Also, the information content of the volatility spread improves significantly when we adjust for skewness. This new stylized fact should help to discriminate among competing theories of the volatility spread.

We develop the practitioner’s version of the HG model, but restrict deviations from a flat volatility-skewness surface to deviations in excess kurtosis. Empirically, models based on the HG distribution perform better than their Gaussian counterpart. Hedging performance are also substantially improved. We introduce the implied volatility and skewness surface, which we study empirically. Finally, we document the term structure of volatility, skewness, and kurtosis out to an horizon of 9 months. We find that dependence in returns have a larger impact on skewness than kurtosis, highlighting again the potential information content of skewness.

Ongoing research evaluates the robustness of the results to the use of non-parametric measures of option-implied moments and to the presence of other known predictors of SP500 excess returns.

A Appendix

A Proposition 1

Our candidate SDF is, for given ν ,

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

where Ψ is the log-cumulant function of ε ,

$$\Psi(u) = 2u \frac{\sqrt{h(\Delta)}}{\alpha(\Delta)} - \frac{4}{\alpha(\Delta)^2} \ln \left[1 + \frac{1}{2} u \alpha(\Delta) \sqrt{h(\Delta)} \right].$$

Following CEFJ, this SDF defines an Equivalent Martingale Measure [EMM] if and only if

$$\Psi(\nu(\Delta) - 1) - \Psi(\nu(\Delta)) - \Psi(-1) + (\mu - r_0) \Delta = 0,$$

which has the following, unique, solution for $\nu(\Delta)$,

$$\nu(\Delta) = -\frac{2}{\alpha(\Delta) \sqrt{h(\Delta)}} + \frac{g(\Delta)}{g(\Delta) - 1},$$

where

$$g(\Delta) = \exp \left(-\frac{(\mu - r_0) \Delta}{4} \alpha(\Delta)^2 + \frac{\alpha(\Delta) \sqrt{h(\Delta)}}{2} \right).$$

Proposition 2 of CEFJ establishes sufficient conditions on Ψ for the solution to be unique.

B Limit of Risk-Neutral Volatility

Define

$$\begin{aligned} \Pi_0 &\equiv (\mu - r) \\ \beta(\Delta) &\equiv \alpha(\Delta) \frac{\sqrt{\Delta}}{2} \\ \sigma^*(\Delta) &\equiv \sqrt{h^*(\Delta)} / \sqrt{\Delta}, \end{aligned}$$

and note that the drift correction term can be written as

$$2 \frac{\sqrt{h^*(\Delta)} - \sqrt{h(\Delta)}}{\alpha(\Delta)} = \frac{\sigma^*(\Delta) - \sigma}{\beta(\Delta)} \Delta. \quad (17)$$

We first study the limit of the numerator as skewness tends to zero. Using the definitions above we have (see Proposition 2)

$$\sigma^*(\Delta) = \frac{g(\beta(\Delta)) - 1}{\beta(\Delta) g(\beta(\Delta))} \quad (18)$$

where, with a slight abuse of notation,

$$g(\beta(\Delta)) \equiv \exp(-\Pi_0 \beta(\Delta)^2 + \beta(\Delta) \sigma), \quad (19)$$

which leads to an indeterminacy when skewness tends to zero. We use the first order expansion of the exponential function, $\exp(x) = 1 + x + x\theta(x)$ where $\theta(x)$ tends to zero when x tends to zero. Substituting in Equation 18 leads to, after some simplification,

$$\sigma^*(\Delta) = \frac{-\Pi_0 \beta(\Delta) + \sigma + \theta(\beta(\Delta))}{1 - \Pi_0 \beta(\Delta)^2 + \beta(\Delta) \sigma + \beta(\Delta) \theta(\beta(\Delta))},$$

and taking the limit shows that $\sigma^*(\Delta) \rightarrow \sigma$ when $\beta(\Delta) \rightarrow 0$.

Note then that the limit of 17 leads to an indeterminacy. We will again apply a Taylor expansion but, first, we compute the first order derivative of (18) with respect to $\beta(\Delta)$ using Equation (19) to compute the derivative of $g(\beta(\Delta))$ which leads to

$$\frac{d\sigma^*(\beta(\Delta))}{d\beta(\Delta)} = \frac{1 - g(\beta(\Delta)) + \beta(\Delta)(\sigma - 2\Pi_0\beta(\Delta))}{\beta(\Delta)^2 g(\beta(\Delta))},$$

where again we face an indeterminacy. We use a second-order expansion of $g(\beta(\Delta))$

$$g(\beta(\Delta)) = g(0) + \beta(\Delta)g'(0) + \frac{1}{2}g''(0)\beta(\Delta)^2 + \beta(\Delta)^2\theta(\beta(\Delta)),$$

where $\theta(\beta(\Delta))$ tends to zero when $\beta(\Delta)$ tends to zero. Substituting these results in a first-order expansion for $\sigma^*(\beta(\Delta))$,

$$\sigma^*(\beta(\Delta)) = \sigma^*(0) + \frac{d\sigma^*(\beta(\Delta))}{d\beta(\Delta)}(0)\beta(\Delta) + \beta(\Delta)\theta(\beta(\Delta)),$$

leads to

$$\frac{\sigma^*(\Delta) - \sigma}{\beta(\Delta)} = -\left(\Pi_0 + \frac{\sigma^2}{2}\right) + \theta(\beta(\Delta)),$$

which, in the limit, delivers the desired result. Note that we then have

$$\mu\Delta + 2\frac{\sqrt{h^*(\Delta)} - \sqrt{h(\Delta)}}{\alpha(\Delta)} = \left(r - \frac{\sigma^2}{2}\right)\Delta + \Delta\theta(\beta(\Delta)).$$

and, finally, that if we substitute the second-order expansion for $g(\Delta)$ in the solution for ν , we get

$$\nu(\Delta) \rightarrow \frac{\mu - r + \frac{\sigma^2}{2}}{\sigma^2} = \frac{\mu - r}{\sigma^2} + \frac{1}{2},$$

C Taylor Expansion of the Price of Risk

We want to show that,

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

where

$$\begin{aligned}\nu(\beta) &= -\frac{1}{\beta\sigma} + \frac{g(\beta)}{g(\beta) - 1} \\ g(\beta) &= \exp(-(\mu - r)\beta^2 + \beta\sigma).\end{aligned}$$

Recall that $\nu(0) = (\mu - r)/\sigma^2 + \frac{1}{2}$ and note that

$$\begin{aligned}\nu'(\beta) &= \frac{1}{\beta^2\sigma} - \frac{g'(\beta)}{(g(\beta) - 1)^2} \\ g'(\beta) &= (-2(\mu - r)\beta + \sigma)g(\beta),\end{aligned}$$

We evaluate the limit of this derivative as $\beta \rightarrow 0$ using, as above, the second-order expansion of $g(\beta)$. After tedious but straightforward computations, the result is

$$\begin{aligned}\nu'(0) &= \frac{(\mu - r)^2 + \frac{\sigma^4}{4} - 2(\mu - r)\sigma^2 + \frac{\sigma^4}{3} - (\mu - r)\sigma^2 + 2(\mu - r)\sigma^2 + (\mu - r)\sigma^2 - \frac{\sigma^4}{2}}{\sigma^3} \\ &= \frac{(\mu - r)^2 + \frac{\sigma^4}{12}}{\sigma^3}.\end{aligned}$$

D Proposition 2

From CEFJ, the logarithm risk-neutral of the risk-neutral Moment Generating Function is

$$\begin{aligned}\Psi^{Q^*}(u) &= -u\Psi'(\nu(\Delta)) + \Psi(\nu(\Delta) + u) - \Psi(\nu(\Delta)) \\ &= 2u\frac{\sqrt{h^*(\Delta)}}{\alpha(\Delta)} - \frac{4}{\alpha(\Delta)^2} \ln \left[1 + \frac{1}{2}u\alpha(\Delta)\sqrt{h^*(\Delta)} \right],\end{aligned}$$

implying that

$$\varepsilon_{t+\Delta}^* = \frac{\sqrt{h(\Delta)}}{\sqrt{h^*(\Delta)}}\varepsilon_{t+\Delta} + \nu(\Delta)\sqrt{h(\Delta)}.$$

The HG model can then be written as

$$\ln(S_{t+\Delta}/S_t) = r_0\Delta - \gamma^*(\Delta) + \sqrt{h^*(\Delta)}\varepsilon_{t+\Delta}^*,$$

where

$$\gamma^*(\Delta) = \Psi^{Q^*}(-1) = -2\frac{\sqrt{h^*(\Delta)}}{\alpha(\Delta)} - \frac{4}{\alpha(\Delta)^2} \ln \left[1 - \frac{1}{2}\alpha(\Delta)\sqrt{h^*(\Delta)} \right],$$

and with

$$\sqrt{h^*(\Delta)} = \frac{2(g(\Delta) - 1)}{\alpha(\Delta)g(\Delta)}.$$

Substituting back in the equation for returns under the risk-neutral measure, and simplifying, yields the results.

E Greeks

For notational simplicity we introduce $a \equiv H/\beta(\Delta)^2$. We begin with the sensitivity to changes in the underlying stock price. The HG option price is homogenous of degree one in stock price and strike. Then the standard result holds and the option delta is simply

$$\frac{\partial C_t}{\partial S_t} = C_{1,t}, \quad (20)$$

which depends on skewness. Next, the sensitivity of the option's delta with respect to the stock price is

$$\frac{\partial^2 C_t}{\partial S_t^2} = \frac{e^{-(d_2+r_f H)} d_2^{a-1} K}{|\beta|\sigma^*\Gamma(a)} \frac{1}{S_t^2}, \quad (21)$$

which also depends on skewness and moneyness. The sensitivity of option prices to changes in the underlying risk-neutral volatility is

$$\frac{\partial C_t}{\partial \sigma_t^*} = \frac{|\beta|\sigma^*e^{(-r_f H)} K}{\sigma^*(1-\beta\sigma^*)} \frac{e^{-d_2} d_2^a}{\Gamma(a)}, \quad (22)$$

and, finally, the sensitivity of option prices to changes in the skewness of returns is given by

$$\begin{aligned}\frac{\partial C_t}{\partial \beta} &= -\frac{2a}{\beta} \left[(\ln(d_2) - \Psi(a)) C_t - K e^{(-r_f H)} P(a, d_2) \ln(1 - \beta\sigma) \right] \\ &+ \frac{2a}{\beta} \Gamma(a) d_2^a S_t (1 - \beta\sigma)^a {}_2\bar{F}_2(a, a; a+1, a+1; -d_1) \\ &- \frac{2a}{\beta} \Gamma(a) d_2^a K e^{(-r_f H)} \bar{F}_2(a, a; a+1, a+1; -d_2) \\ &- K e^{(-r_f H)} \frac{\sigma^*}{1 - \beta\sigma^*} \frac{e^{-d_2} d_2^a}{\Gamma(a)},\end{aligned} \quad (23)$$

where

$$\Psi(a, z) = P(a, z) \ln(z) - \Gamma(a) z^a {}_2\bar{F}_2(a, a; a+1, a+1; -z),$$

and where ${}_2\bar{F}_2(\cdot)$ is the regularized hypergeometric function.

F Proposition 3

A no-arbitrage price of a European call option with strike price K and maturity T can be obtained from the computation of the discounted expectation of the terminal payoff under the risk-neutral measure. That is,

$$\begin{aligned} C_t(K, M) &= E^Q [\max(S_{t+T} - K, 0)] \\ C_t &= \exp(-rT) S_t E^Q [\exp(R_{t,M}) 1_{[R_{t,M} > \ln(K/S_t)]}] - \exp(-r_0 T) K P^Q [R_{t,M} > \ln(K/S_t)]. \end{aligned}$$

We can compute $P^Q[R_{t,M} > \ln(K/S_t)]$ from the distribution function of a gamma variable. Note first that

$$P^Q[R_{t,M} > \ln(K/S_t)] = P^Q \left[\frac{\beta(\Delta)}{\sqrt{\Delta M}} y_{t,M}^* > \frac{\ln(K/S_t) - \mu^*(\Delta) M \Delta}{\sqrt{\Delta M} \sigma^*(\Delta)} + \frac{\sqrt{\Delta M}}{\beta(\Delta)} \right],$$

where we define

$$\frac{2\sqrt{M}}{\alpha(\Delta)} \left(\varepsilon_{t,M}^* + \frac{2\sqrt{M}}{\alpha(\Delta)} \right) = y_{t,M}^* \sim^Q \Gamma \left(\frac{4M}{\alpha(\Delta)^2}, 1 \right),$$

based on the characterization of the standardized Gamma distribution given in Equation 2. If $\alpha(\Delta) > 0$,

$$P^Q[R_{t,M} > \ln(K/S_t)] = \frac{\Gamma \left(\frac{T}{\beta^2(\Delta)}; \frac{T}{\beta^2(\Delta)} + \frac{\ln(K/S_t) - \mu^*(\Delta) T}{\beta(\Delta) \sigma^*(\Delta)} \right)}{\Gamma \left(\frac{T}{\beta^2(\Delta)} \right)},$$

where $\Gamma(a, x)$ is the upper incomplete gamma function²⁵ and if $\alpha(\Delta) < 0$,

$$\begin{aligned} P^Q[R_{t,M} > \ln(K/S_t)] &= \frac{\gamma \left(\frac{T}{\beta(\Delta)^2}; \frac{T}{\beta(\Delta)^2} + \frac{\ln(K/S_t) - \mu^* T}{\beta(\Delta) \sigma^*(\Delta)} \right)}{\Gamma \left(\frac{T}{\beta(\Delta)^2} \right)} \\ &= 1 - \frac{\Gamma \left(\frac{T}{\beta(\Delta)^2}; \frac{T}{\beta(\Delta)^2} + \frac{\ln(K/S_t) - \mu^* T}{\beta(\Delta) \sigma^*(\Delta)} \right)}{\Gamma \left(\frac{T}{\beta(\Delta)^2} \right)}. \end{aligned}$$

Similarly,

$$\begin{aligned} &E^Q [\exp(R_{t,M}) 1_{[R_{t,M} > \ln(K/S_t)]}] \\ &= \exp \left(\mu^*(\Delta) M \Delta - \frac{\sigma^*(\Delta) M \Delta}{\beta(\Delta)} \right) E^Q \left[\exp(\sigma^*(\Delta) \beta(\Delta) y_{t,M}^*) 1_{\left[\frac{\beta(\Delta)}{\sqrt{\Delta M}} y_{t,M}^* > \kappa \right]} \right] \end{aligned}$$

where we use

$$\kappa = \frac{\ln(K/S_t) - \mu^*(\Delta) M \Delta}{\sqrt{\Delta M} \sigma^*(\Delta)} + \frac{\sqrt{\Delta M}}{\beta(\Delta)}.$$

Then, if $\alpha(\Delta) > 0$, and using that $y_{t,M}^*$ has a standard gamma distribution with parameter $\frac{M\Delta}{\beta(\Delta)^2}$, we have

$$\begin{aligned} &E^Q \left[\exp(\sigma^*(\Delta) \beta(\Delta) y_{t,M}^*) 1_{\left[y_{t,M}^* > \frac{\sqrt{\Delta M} \kappa}{\beta(\Delta)} \right]} \right] \\ &= \int_{(1 - \sigma^*(\Delta) \beta(\Delta)) \frac{\sqrt{\Delta M} \kappa}{\beta(\Delta)}}^{\infty} \exp(-z_{t,M}^*) \frac{(z_{t,M}^*)^{\frac{M\Delta}{\beta(\Delta)^2} - 1}}{(1 - \sigma^*(\Delta) \beta(\Delta))^{\frac{M\Delta}{\beta(\Delta)^2}} \Gamma \left(\frac{M\Delta}{\beta(\Delta)^2} \right)} dz_{t,M}^* \\ &= \frac{\Gamma \left(\frac{M\Delta}{\beta(\Delta)^2}; \left(\frac{M\Delta}{\beta(\Delta)^2} + \frac{\ln(K/S_t) - \mu^*(\Delta) \Delta M}{\beta(\Delta) \sigma^*(\Delta)} \right) (1 - \sigma^*(\Delta) \beta(\Delta)) \right)}{\Gamma \left(\frac{M\Delta}{\beta(\Delta)^2} \right) (1 - \sigma^*(\Delta) \beta(\Delta))^{\frac{M\Delta}{\beta(\Delta)^2}}} \end{aligned}$$

²⁵The upper incomplete gamma function is defined as $\Gamma(a, x) = \int_x^{\infty} t^{a-1} e^{-t} dt$ while the lower incomplete gamma function is defined as $\gamma(a, x) = \int_0^x t^{a-1} e^{-t} dt$. Note that $\Gamma(a) = \Gamma(a, 0)$ while $\gamma(a) = \gamma(a, \infty)$.

and, using the change of variables $(1 - \sigma^*(\Delta) \beta(\Delta)) y_{t,M}^* = z_{t,M}^*$, it follows that

$$\begin{aligned} & E^Q \left[\exp(R_{t,M}) 1_{[R_{t,M} > \ln(K/S_t)]} \right] \\ &= \exp \left(\left(\mu^*(\Delta) - \frac{\sigma^*(\Delta)}{\beta(\Delta)} \right) T \right) \frac{\Gamma \left(\frac{T}{\beta(\Delta)^2}; \left(\frac{T}{\beta(\Delta)^2} + \frac{\ln(K/S_t) - \mu^*(\Delta)T}{\beta(\Delta)\sigma^*(\Delta)} \right) (1 - \sigma^*(\Delta) \beta(\Delta)) \right)}{\Gamma \left(\frac{T}{\beta(\Delta)^2} \right) (1 - \sigma^*(\Delta) \beta(\Delta))^{\frac{T}{\beta(\Delta)^2}}}. \end{aligned}$$

If, however, $\alpha(\Delta) < 0$ then

$$\begin{aligned} & E^Q \left[\exp(\sigma^*(\Delta) \beta(\Delta) y_{t,M}^*) 1_{\left[\frac{\beta(\Delta)}{\sqrt{\Delta M}} y_{t,M}^* > \kappa \right]} \right] \\ &= \frac{\gamma \left(\frac{M\Delta}{\beta(\Delta)^2}; \left(\frac{M\Delta}{\beta(\Delta)^2} + \frac{\ln(K/S_t) - \mu^*(\Delta)\Delta M}{\beta(\Delta)\sigma^*(\Delta)} \right) (1 - \sigma^*(\Delta) \beta(\Delta)) \right)}{\Gamma \left(\frac{M\Delta}{\beta(\Delta)^2} \right) (1 - \sigma^*(\Delta) \beta(\Delta))^{\frac{M\Delta}{\beta(\Delta)^2}}}, \end{aligned}$$

and then

$$\begin{aligned} & E^Q \left[\exp(R_{t,M}) 1_{[R_{t,M} > \ln(K/S_t)]} \right] \\ &= \exp \left(\left(\mu^*(\Delta) - \frac{\sigma^*(\Delta)}{\beta(\Delta)} \right) T \right) \frac{\gamma \left(\frac{T}{\beta(\Delta)^2}; \left(\frac{T}{\beta(\Delta)^2} + \frac{\ln(K/S_t) - \mu^*(\Delta)T}{\beta(\Delta)\sigma^*(\Delta)} \right) (1 - \sigma^*(\Delta) \beta(\Delta)) \right)}{\Gamma \left(\frac{T}{\beta(\Delta)^2} \right) (1 - \sigma^*(\Delta) \beta(\Delta))^{\frac{T}{\beta(\Delta)^2}}}. \end{aligned}$$

G Proposition 4

Suppose that the underlying stock price evolution under the risk-neutral measure is given by

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

where δ is a risk-adjustment factor, y is a random number with mean zero, variance 1, skewness, $\frac{2\beta}{\sqrt{T}}$ and kurtosis, λ_2 . Suppose also that the probability density of y is described by the following Edgeworth series expansion around the standardized gamma distribution:

$$f(y) = g(y) + \frac{\lambda_2 - \frac{6\beta^2}{T}}{4!} \frac{d^4 g(y)}{dy^4},$$

where $g(y)$ is the standardized gamma density function given by

$$g(y) = \frac{\sqrt{T}z^{a-1}e^{-z}}{|\beta|\Gamma(a)} \quad \text{if } \beta y > -\sqrt{T},$$

and where $z = \frac{\sqrt{T}}{\beta}y + a$. Imposing that gross stock returns are a martingale under the risk-neutral measure,

$$\begin{aligned} E_0^Q[\exp(R_T)] &= E_0^Q[\exp((r - \delta)T + \sigma\sqrt{T}y)] \\ &= \exp((r - \delta)T) \int \exp(\sigma\sqrt{T}y) \left[g(y) + \frac{\lambda_2 - \frac{6\beta^2}{T}}{4!} \frac{d^4 g(y)}{dy^4} \right] dy, \end{aligned}$$

leads to the required risk-adjustment,

$$\delta = \frac{1}{T} \ln \left[\exp \left(-\frac{\sigma T}{\beta} - a \ln(1 - \beta\sigma) \right) + \frac{\lambda_2 - \frac{6\beta^2}{T}}{4!} \int \exp(\sigma\sqrt{T}y) \frac{d^4 g(y)}{dy^4} dy \right].$$

The price of a European call option is

$$c_0^* = e^{-rT} \int_{-d_2^*}^{\infty} \left(S_0 \exp \left((r - \delta) T + \sigma \sqrt{T} y \right) - K \right) f(y) dy$$

where

$$d_2^* = \frac{\ln(S_0/K) + (r - \delta) T}{\sigma \sqrt{T}}.$$

We have

$$c_0^* = e^{-rT} \left[S_0 \exp((r - \delta) T) \int_{-d_2^*}^{\infty} \exp(\sigma \sqrt{T} y) f(y) dy - K \int_{-d_2^*}^{\infty} f(y) dy \right].$$

For the first integral, we have

$$\int_{-d_2^*}^{\infty} \exp(\sigma \sqrt{T} y) f(y) dy = \int_{-d_2^*}^{\infty} \exp(\sigma \sqrt{T} y) g(y) dy + \kappa \int_{-d_2^*}^{\infty} \exp(\sigma \sqrt{T} y) \frac{d^4 g(y)}{dy^4} dy$$

and for $\beta \leq 0$, say, and $d_1^* = \bar{d}_2(1 - \sigma\beta)$ we have

$$\int_{-d_2^*}^{\infty} \exp(\sigma \sqrt{T} y) g(y) dy = \frac{e^{-\sigma\beta a}}{(1 - \sigma\beta)^a} P(a, d_1^*)$$

while

$$\begin{aligned} & \int_{-d_2^*}^{\infty} \exp(\sigma \sqrt{T} y) \frac{d^4 g(y)}{dy^4} dy \\ &= a^2 e^{-\sigma\beta a} \left[\frac{P(a-4, d_1^*)}{(1-\sigma\beta)^{a-4}} - 4 \frac{P(a-3, d_1^*)}{(1-\sigma\beta)^{a-3}} \right. \\ & \quad \left. + 6 \frac{P(a-2, d_1^*)}{(1-\sigma\beta)^{a-2}} - 4 \frac{P(a-1, d_1^*)}{(1-\sigma\beta)^{a-1}} + \frac{P(a, d_1^*)}{(1-\sigma\beta)^a} \right]. \end{aligned}$$

Next, for the second integral above,

$$\int_{-d_2^*}^{\infty} f(y) dy = \int_{-d_2^*}^{\infty} g(y) dy + \kappa \int_{-d_2^*}^{\infty} \frac{d^4 g(y)}{dy^4} dy$$

with

$$\begin{aligned} \int_{-d_2^*}^{\infty} g(y) dy &= \int_0^{a-d_2^* \frac{\sqrt{T}}{\beta}} \frac{z^{a-1} e^{-z}}{\Gamma(a)} dz = P(a, \bar{d}_2) \\ \int_{-d_2^*}^{\infty} \frac{d^4 g(y)}{dy^4} dy &= a^2 \left[\frac{P(a-4, \bar{d}_2) - 4P(a-3, \bar{d}_2) + 6P(a-2, \bar{d}_2) - 4P(a-1, \bar{d}_2) + P(a, \bar{d}_2)}{6P(a-2, \bar{d}_2) - 4P(a-1, \bar{d}_2) + P(a, \bar{d}_2)} \right]. \end{aligned}$$

H Identifying Restriction on the P-HG

The equality of prices from the true model and the P-HG for at-the-money options implies that

$$\begin{aligned} P(a, d_1^*) - P(a, d_2^*) &= P(a, d_1) - (1 + T^2 \sigma^4 \kappa) P(a, d_2) \\ & \quad + \kappa \frac{T^2 \sigma}{\beta^3} [-h''(d_2) + \sigma \beta h'(d_2) - \sigma^2 \beta^2 h(d_2)], \end{aligned}$$

while the equality of the first derivative of prices implies

$$\begin{aligned} P(a, d_2^*) + \frac{\sigma_{I0} \gamma_1}{\bar{\sigma} \sqrt{T}} \frac{d_2^* \beta h(d_2^*)}{(1 - \beta \sigma_{0I})} &= (1 + T^2 \sigma^4 \kappa) P(a, d_2) \\ & \quad + \kappa a^2 [h'''(d_2) + \sigma^3 \beta^3 h(d_2)], \end{aligned}$$

and, finally, the equality of the second derivatives implies

$$\frac{h(d_2^*)}{\sigma_{0I}} \left[1 + \frac{(2a - d_1^* - d_2^*) \beta \sigma_{I0} \gamma_1}{(1 - \beta \sigma_{0I}) \bar{\sigma} \sqrt{T}} + \frac{d_2^* \beta^3 \sigma_{I0}^3 \gamma_1^2}{(1 - \beta \sigma_{0I})^2 \bar{\sigma}^2 T} + \frac{2d_2^* \beta^2 \sigma_{I0}^2 \gamma_2}{(1 - \beta \sigma_{0I}) \bar{\sigma}^2 T} \right] =$$

$$(1 + T^2 \sigma^4 \kappa) \frac{h(d_2)}{\sigma} + \frac{\kappa a^2}{\sigma} \left[h^{(4)}(d_2) + \sigma^3 \beta^3 h'(d_2) \right].$$

Then, linearizing the left sides of the equations around $\sigma_{I0} = \sigma$, $\gamma_1 = 0$ and $\gamma_2 = 0$, respectively, and the right side around $\kappa = 0$ leads to

$$\frac{\sigma_{I0} - \sigma}{\sigma} = a^2 (1 - \sigma \beta) \left(\sigma^3 \beta^3 \frac{P(a, d_2)}{h(d_2) d_2} + \frac{(a-1)(a-2)}{d_2^3} - \frac{(a-1)(2 + \sigma \beta)}{d_2^2} + \frac{1 + \sigma \beta + \sigma^2 \beta^2}{d_2} \right) \kappa$$

$$\gamma_1 = -\frac{\bar{\sigma} \sqrt{T} (a - d_1)}{\beta \sigma_2} \frac{\sigma_{I0} - \sigma}{\sigma} + \frac{\bar{\sigma} \sqrt{T} a^2 (1 - \sigma \beta)}{d_2} \left[\frac{\beta^3 \sigma^3 P(a, d_2)}{h(d_2)} + 2\beta^2 \sigma^2 + \frac{h^{(3)}(d_2)}{\beta \sigma h(d_2)} \right] \kappa$$

$$\gamma_2 = -\frac{\bar{\sigma}^2 T}{2\beta^2 \sigma^2 d_2} \left(\frac{h'(d_2)(a - d_1)}{h(d_2)} - 1 + \sigma \beta \right) \frac{\sigma_{I0} - \sigma}{\sigma}$$

$$+ -\frac{\bar{\sigma} \sqrt{T} (2a - d_1 - d_2)}{2d_2 \beta \sigma} \gamma_1 + \left(\sigma + \frac{h'(d_2)}{\beta h(d_2)} \right) \frac{\sigma (1 - \sigma \beta) \bar{\sigma}^2 T^3}{2d_2 \beta^2} \kappa$$

where

$$d_2 = \frac{-a \ln(1 - \sigma \beta)}{\sigma \beta}$$

$$d_1 = d_2 (1 - \sigma \beta)$$

$$a = \frac{T}{\beta^2}$$

$$\kappa = \frac{\lambda_2 - \frac{6\beta^2}{T}}{4!}.$$

References

- Alexander, C. and L. M. Nogueira (2005). Model-free hedge ratios and scale-invariant models. *Journal of Banking and Finance* 31.
- 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.
- Bakshi, G. and N. Kapadia (2003). Delta-hedged gains and the negative market volatility risk premium. *Review of Financial Studies* 16.
- 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.
- Bakshi, G. and D. Madan (2000). Spanning and derivative-security valuation. *Journal of Financial Economics* 58.
- Bakshi, G. and D. Madan (2006). A theory of volatility spreads. *Management Science* 52.
- Bates, D. (2000). Post-'87 crash fears in the sp500 futures option market. *Journal of Econometrics* 94.
- Bates, D. (2005). Hedging the smirk. *Financial Research Letters* 2.
- 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 (2008). Variance risk premiums. *Review of Financial Studies* 000.
- Christoffersen, P., R. Elkamhi, B. Feunou, and K. Jacobs (2007). Option valuation with conditional heteroskedasticity and non-normality. *McGill University*.
- Christoffersen, P., S. Heston, and K. Jacobs (2006). Option valuation with conditional skewness. *Journal of Econometrics* 131.
- Christoffersen, P., K. Jacobs, and G. Vainberg (2008). Forward-looking betas. *Unpublished*.

- 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.
- Dumas, B., J. Fleming, and R. Whaley (1998). Implied volatility functions: Empirical tests. *The Journal of Finance* 53.
- Galai, D. (1983). The components of the returns from hedging options against stocks. *The Journal of Business* 56.
- 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. working paper.
- 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.
- Zhang, J. and Y. Xiang (2005). Implied volatility smirk. Working Paper.

Table I: Summary statistics for strike price and maturity categories.

(a) Summary statistics by moneyness

	Moneyness					All
	<0.95	0.95 to 0.975	0.975 to 1	1 to 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. For each moneyness and strike price category, the first line gives the number of contracts and the second line give the average Implied Volatility.

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

Table II: **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 . The sample period is from January 1996 to December 2004.

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

Table III: Weekly in-sample 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 β 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.

(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: 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 β 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.

(a) Out-of-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	0.795	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					
	1	2	3	4-6	7-9	All
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: 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 β 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.

(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: 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 β 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.

(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 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 β 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.


(a) Monthly In-Sample RMSE

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
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

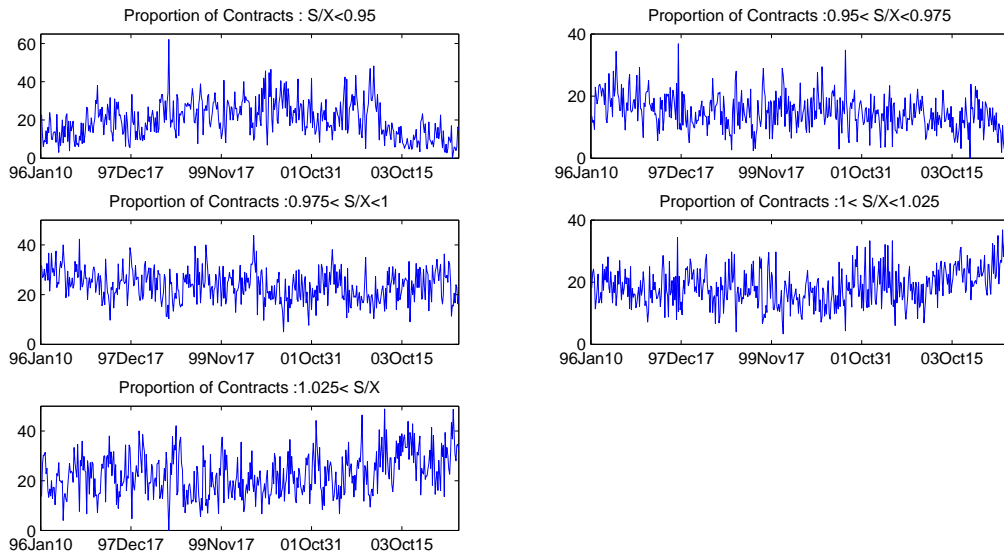
(b) Monthly Out-of-Sample RMSE

Model	Moneyness					
	$S/X < 0.95$	$0.95 < S/X < 0.975$	$0.975 < S/X < 1$	$1 < S/X < 1.025$	$1.025 < S/X$	All
HG	0.908	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

Figure 1: Number of call option contracts at each date




(a) Total number of contracts.




(b) Proportion of contracts in each moneyness category.

Figure 2: Time series of implied volatility and implied skewness from the smoothed version of the SP-HG2 model. This is a practitioner's version of the Homoscedastic Gamma model where the IV curve is restricted to depends only on the (constant) excess kurtosis.




(a) Implied Volatility




(b) Implied Skewness

Figure 3: 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 $\frac{\ln(S/K)(-r\tau)}{\bar{\sigma}\sqrt{\tau}}$ to correct for maturity differences.




(a) Level




(b) Percentage Deviation from BSM


Figure 4: Implied volatility and skewness surfaces for different maturity categories where moneyness is defined as $\ln(S/K)(-r\tau)$. Maturity groups are defined using settlement dates.




(a) Month 1




(b) Month 2



(c) Month 3




(d) Months 4 to 6




(e) Months 7 to 9


Figure 5: 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.




(a) Month 1




(b) Month 2



(c) Month 3



(d) Months 4 to 6



(e) Months 7 to 9

Figure 6: 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.

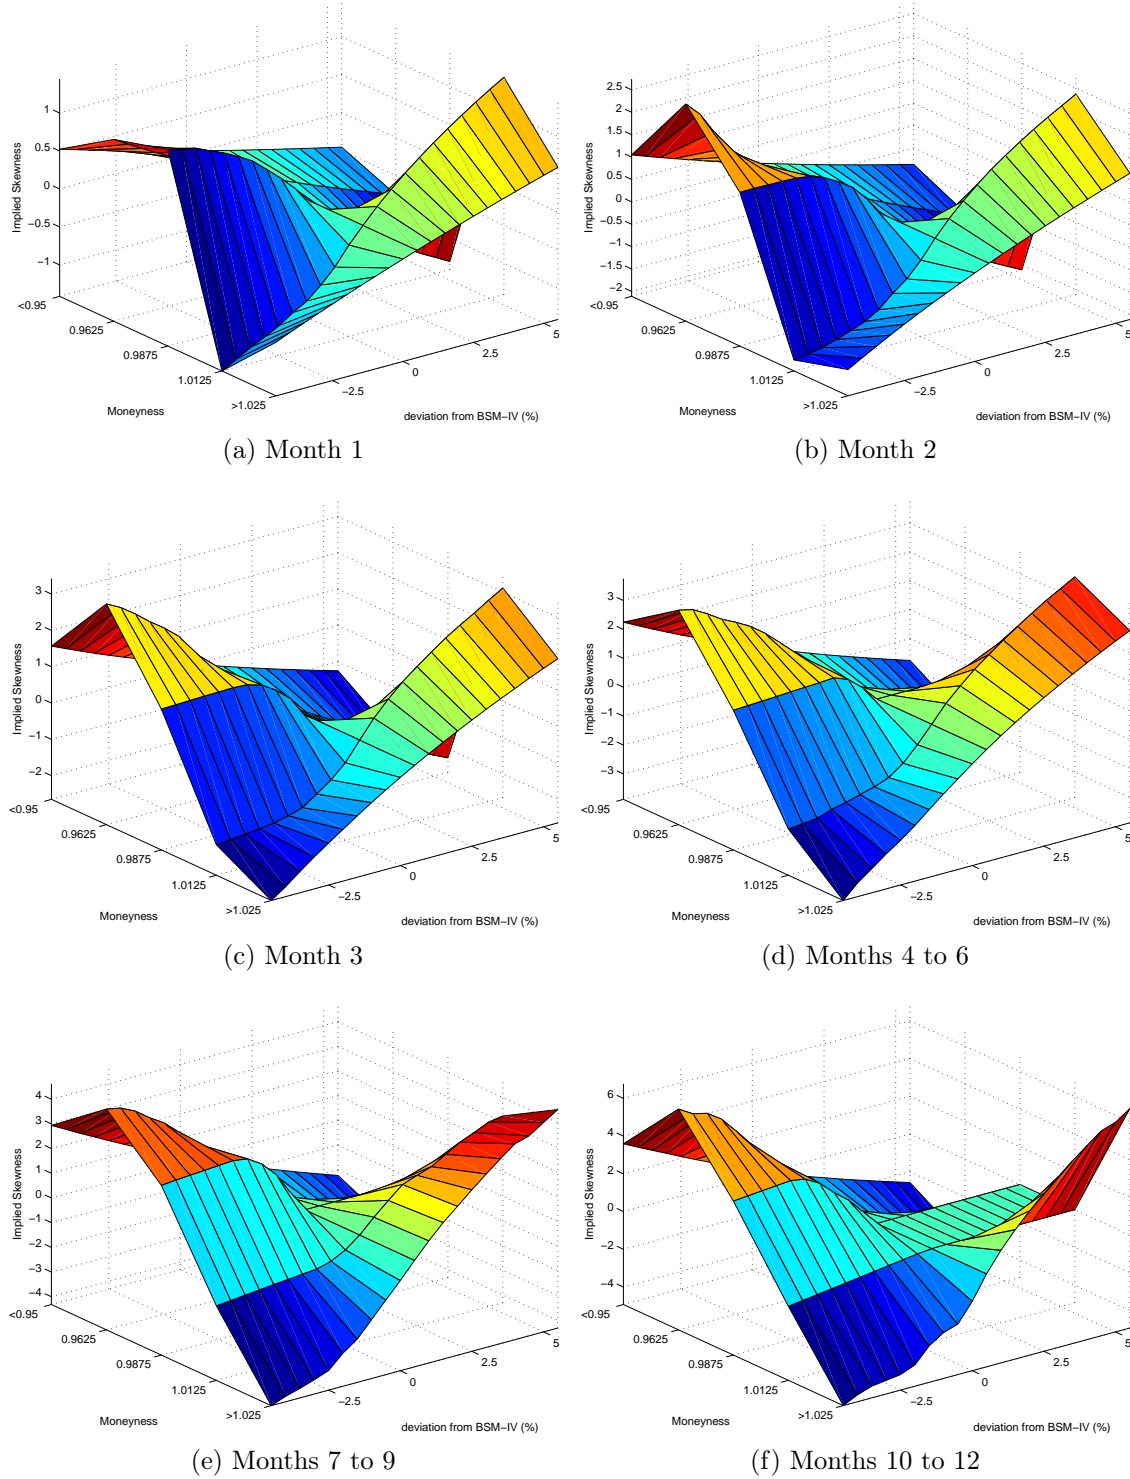





Figure 7: Time-series average of estimates of $\Theta_t = (\sigma_t, \gamma_{1,t}, \gamma_{2,t})$ from the P-HG3 (unrestricted) model but for different values of skewness. The parameters govern the IV curve: $\sigma_{i,t} = \sigma_{I0,t}(1 + \gamma_{1,t}\xi_i + \gamma_{2,t}\xi_i^2)$.



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




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




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


Figure 8: 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.



(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.


Figure 9: Term Structure of implied volatility, (minus) the implied skewness and (minus) the implied excess kurtosis from the SP-HG2 model.



(a) Implied Volatility



(b) Implied Skewness



(c) Implied Kurtosis