

CREATES Research Paper 2008-13

Option Pricing using Realized Volatility

Lars Stentoft



School of Economics and Management
University of Aarhus
Building 1322, DK-8000 Aarhus C
Denmark



Option Pricing using Realized Volatility*

Lars Stentoft[†]

HEC Montréal and CREATES

February 29, 2008

Abstract

In the present paper we suggest to model Realized Volatility, an estimate of daily volatility based on high frequency data, as an Inverse Gaussian distributed variable with time varying mean, and we examine the joint properties of Realized Volatility and asset returns. We derive the appropriate dynamics to be used for option pricing purposes in this framework, and we show that our model explains some of the mispricings found when using traditional option pricing models based on interdaily data. We then show explicitly that a Generalized Autoregressive Conditional Heteroskedastic model with Normal Inverse Gaussian distributed innovations is the corresponding benchmark model when only daily data is used. Finally, we perform an empirical analysis using stock options for three large American companies, and we show that in all cases our model performs significantly better than the corresponding benchmark model estimated on return data alone. Hence the paper provides evidence on the value of using high frequency data for option pricing purposes.

JEL Classification: C22, C53, G13

Keywords: Option Pricing, Realized Volatility, Stochastic Volatility, GARCH.

*The author thanks Timothy Simmons, Phelim Boyle, and Éric Jacquier, participants at the CREATES opening conference and the 2007 Northern Finance Association annual meeting, as well as seminar participants at HEC Montreal for valuable comments. Financial support from CREATES, funded by the Danish National Research Foundation, and from IFM2 is gratefully appreciated.

[†]Email: lars.stentoft@hec.ca. Department of Finance, HEC Montréal, 3000 chemin de la Côte-Sainte-Catherine, Montréal (Québec), Canada, H3T 2A7.

1 Introduction

In many finance applications it is of paramount importance to measure volatility. This is particularly so when the ultimate goal is the pricing of derivatives securities like options where volatility is one of the important drivers of the price. For this reason the modelling and forecasting of volatility has been the focus of much work. While volatility initially was treated as being constant through time, e.g. as it was the underlying assumption in the models of Black & Scholes (1973) and Merton (1973), it is by now well understood that the volatility of most financial return series varies through time. Many of the extensions to the classical constant volatility framework which allow for time varying volatility fall within either the discrete time Generalized Autoregressive Conditional Heteroskedastic (GARCH) framework or within the continuous time Stochastic Volatility (SV) framework.

The GARCH framework which was initially proposed by Engle (1982) and extended by Bollerslev (1986) has been applied extensively in the empirical finance literature (see e.g. the surveys of Bollerslev, Chou & Kroner (1992) and Poon & Granger (2003)). In this framework volatility is treated as a time varying process depending on lagged values of volatility itself and on lagged squared innovations to the returns. Hence in this setting volatility can be estimated entirely from the return data by using the return innovations as a proxy for volatility. In the continuous time SV framework volatility is in fact treated as being “truly” stochastic. Models of this type have been extensively used in the theoretical derivative pricing literature as it is often possible to derive option pricing models more elegantly in a continuous time world. Classical examples include among others the work of Heston (1993), Hull & White (1987), Johnson & Shanno (1987), Scott (1987) and Wiggins (1987). However, when it comes to empirically applying these models, only return data has been available, and hence volatility has in general been treated as an unobserved state variable. In terms of both model estimation and volatility forecasting this complicates the application for empirical option pricing.

Recently the use of high frequency intraday data for estimating daily volatility has received extensive attention with very promising results when compared to using lower frequency interday data. One basic idea is to use sums of squared returns at high frequency, say on an intradaily basis, to estimate volatility at some lower frequency, that is on a daily basis. This idea was formalized in Andersen, Bollerslev, Diebold & Labys (2001) where intraday data is used to calculate what has become known as the Realized Volatility (RV). Since then this area has received much attention and by now various estimators of interday volatility based on intraday information exist (see e.g. Zhang, Mykland & Ait-Sahalia (2005) and Barndorff-Nielsen, Hansen, Lunde & Shephard (2006)). Initially RV measures were used in evaluating volatility forecasts from various time series models like those from GARCH or SV models. A classical reference for this type of exercise is Hansen & Lunde (2005a) where the predictive ability of 330 models within the GARCH framework is compared using this measure. However the potential usefulness of the RV measure is much wider than this,

and it has been used to forecast volatility and returns for e.g. exchange rates in Andersen, Bollerslev, Diebold & Labys (2003). In this paper we analyze the potential value of using RV type measures of daily volatility in an option pricing forecasting context.

In the first part of the paper we analyze three RV series and suggest that the Inverse Gaussian (IG) distribution is an appropriate distribution. However, in order to take account of the serial dependency a more elaborate model is needed and we propose a model for the RV series in which the conditional level of the variance is allowed to be time varying. We refer to this model as the IG-SV(p,q) model, and we show that it is able to capture the time variation in the RV series and provide a good fit of the data. Next we examine the joint properties of the asset return process and the volatility process. In particular, the properties of the standardized returns, that is the returns scaled by the square root of the RV, are examined. The analysis shows that the standardized return series are close to being Gaussian which implies that returns are unconditionally Normal Inverse Gaussian, or NIG, distributed.

In the second part of the paper we derive the risk neutral dynamics to be used for option pricing purposes using the suggested model, and we provide details on how to implement the model. To our knowledge this is the first paper to provide such a modelling framework. We perform a Monte Carlo study to analyze the properties of the model and compare the performance to the well known Black-Scholes-Merton model which is a special case of the suggested framework. This exercise shows that the RV based option pricing model may be able to explain a number of the empirical shortcomings of other option pricing models. In particular, we find that a model using the RV mitigates the underpricing of short term deep out of the money options often found when applying constant volatility models.

In the third part of the paper we relate the proposed model to alternative models based on interday data only. In particular, since the model based on RV data implies that daily returns are NIG distributed with time varying volatility we compare to models within the NIG GARCH framework. We provide estimation results and results on the option pricing properties of these models. Finally, we suggest that the GARCH models may be interpreted as multiplicative Errors in Variables models of the RV based models when RV is unobservable and squared returns are used as a proxy. We justify this interpretation through a Monte Carlo study and by estimating the GARCH model augmented with the RV data. Hence by comparing these models it is possible to gauge the value of using high frequency based models for option pricing purposes.

In the final part of this paper the RV option pricing model is taken to the data in an empirical examination of the pricing performance for options traded on the three individual American stocks. This analysis indicates that our model is in fact able to improve on the mispricing by the classical constant volatility Black-Scholes-Merton framework. However, more importantly the empirical results also show that the RV option pricing model performs significantly better than corresponding models within the NIG GARCH framework which are estimated on return data alone and hence do not use the incremental information available from the high

frequency estimates of volatility. Thus, the results allows us to conclude that there is definite gains to using the RV series calculated from high frequency data in an option pricing context.

The rest of the paper is organized as follows: In Section 2 we describe how the RV estimate of daily variance is calculated and discuss the empirical regularities. This section also discuss the appropriate distribution to be used for modelling the RV process as well as how to model the serial dependency of the time series. In Section 3 we combine the information on the RV with the return process and suggest a framework for modelling both processes jointly. In Section 4 we derive an option pricing model based on and in line with the reported empirical findings and we report results from an extensive Monte Carlo study of the model's pricing properties. In Section 5 we compare the results in terms of estimation and option pricing to existing models based on interday data only and we relate the models to the RV based framework. In Section 6 we report the result from an empirical examination of the pricing performance of the RV model. Finally, in Section 7 we conclude. All tables and figures can be found in the Appendix.

2 Realized Volatility

In order to construct reliable measures of daily volatility it is necessary to obtain data on the underlying assets at high frequency. A well known source of this data is the Trades and Quote Database (TAQ) covering assets traded on the New York Stock Exchange (NYSE). In this work we choose a different database, as we construct a high frequency database from the Berkeley Options DataBase (BODB). In particular, from the BODB it is possible to extract not only relevant option prices but also a series of quotations on the underlying asset as this is reported simultaneously with any quote or trade of the derivatives during the opening hours of the Chicago Board of Options Exchange (CBOE). Thus, with this data we can construct a high frequency data set based on the information that was immediately available to the option traders.

In order to calculate the estimates of daily volatility we follow the procedure outlined in Andersen et al. (2001) which involves calculating the RV as the sum of squared intradaily returns. In particular, we choose to work with intervals of 30 minutes and we pick the observation at or immediately before this interval. Since the options we consider are highly liquid at the CBOE the time between the actual 30 minute points and the time registered for the observation we use is generally small. Based on the 30 minute returns the RV is calculated as the sum of the squared returns over the intraday intervals. That is we calculate

$$RV_t^* = \sum_{i=1}^k R_{t,i}^2, \quad (1)$$

where $R_{t,i}$ is the continuously compounded return at day t between interval $i - 1$ and i and k is the number of intervals during the opening hours of the CBOE which were from 9.00 to 15.00. On a few occasions the estimated variance turns out to be zero. This most likely indicates days with no trades and these observations

are therefore deleted.

To be precise the estimate in (1) is for the variance during the fraction of the day where the exchange is open and thus should not be directly compared to existing measures of daily variance. However, as argued in Hansen & Lunde (2005*b*) there is nothing to suggest that RV_t^* should merely be scaled by the inverse of the fraction of the day considered. The reason is that it is not obvious that an hour with open market should be weighted equally to an hour where the market is closed for volatility calculations. Here we follow the procedure used in Hansen & Lunde (2005*b*) of scaling the Realized Volatility such that the sample average equals that of the squared interday returns, R_t^2 . In particular this scaling is done so that the mean of the two measures of interday variance are equal irrespective of what sample period is considered.

In the booming literature on high frequency data many other suggestions can be found for the construction of the RV series and the potential scaling of the series to a daily level. While more complicated methods could be implemented the procedure used in the present paper provides a valid first approach. In particular, this is the simplest possible procedure and it is easy to replicate. We therefore leave the analysis of the potential improvements on the present results for option pricing using the RV estimated with other more advanced procedures for future research.

2.1 Empirical regularities for Realized Volatility

In the left hand panels of Figure 1 time plots of $\sqrt{RV_t}$ for General Motors (GM), International Business Machines (IBM), and Merck & Company Inc (MRK) are shown. The period used is 1988 through 1995 for a total of 2014 observations for GM, 2016 observations for IBM, and 2013 observations for MRK. Qualitatively, the patterns that arise are quite similar across the three series. In particular, all plots indicate that our estimate of volatility is indeed not constant through time and periods of high RV are followed by low RV periods and vice versa. This confirms the expectations we may have from the large amount of literature documenting time varying volatility patterns for financial data.

In the right hand panels of the figure the autocorrelations for the first 100 lags of $\sqrt{RV_t}$ are plotted. From these plots it is furthermore clear that there is a high degree of dependence in the constructed measure. This finding confirms another popular expectation we may have from the large amount of literature documenting high degrees of dependency in the volatility of financial data.

2.2 Distributional assumptions for Realized Volatility

Previously it has been suggested that the variance of financial data may be well approximated by the lognormal distribution (see e.g. Clark (1973)). However, more recently the Inverse Gaussian (IG) distribution has been suggested as an alternative to the lognormal distribution. In particular, Barndorff-Nielsen & Shephard (2002) argued that the unconditional distribution of the RV process may be well approximated

by the IG distribution. Likewise, in Forsberg & Bollerslev (2002) it is argued that the IG distribution is a suitable distribution for the volatility process of the ECU/USD exchange rate.

The IG density can be written as

$$f_{IG}(z; \sigma^2, a) = \frac{\left(\frac{1}{a\sigma^2}\right)^{-\frac{1}{2}} z^{-\frac{3}{2}}}{(2\pi)^{\frac{1}{2}}} \exp\left(a - \frac{1}{2}\left(\frac{a\sigma^2}{z} + \frac{az}{\sigma^2}\right)\right), \quad (2)$$

and the first four central moments of the distribution are

$$E(z) = \sigma^2, \quad (3)$$

$$V(z) = \frac{\sigma^6}{a}, \quad (4)$$

$$S(z) = 3\sqrt{\frac{\sigma^2}{a}}, \text{ and} \quad (5)$$

$$K(z) = 3 + \frac{15\sigma^2}{a}. \quad (6)$$

We note that if $z \sim IG(\sigma^2, a)$ then $cz \sim IG(c\sigma^2, a)$. Thus, it may be observed that the second parameter does not change under scaling and one may standardize an IG distributed variable by simply scaling it with the mean σ^2 . In Figure 2 we plot $f_{IG}(z; 1, a)$, the standardized density of an IG distributed random variable, for $a = 0.5$, $a = 3$, and $a = 10$ respectively. It is clear from this plot that the IG distribution is quite flexible.

The left hand plots in Figure 3 show the histogram of RV_t together with the density from an appropriately fitted IG distribution.¹ These plots clearly show that the IG distribution is well suited to fit the distribution of RV for the data considered. In Panel A of Table 1 the estimated a parameter is reported together with basic statistics for misspecification. Although the parameter is estimated with appropriate size and quite precisely in this model, the various test statistics for misspecification are significant for all the RV series. This is particularly so for the $Q(20)$ statistics, which indicates the presence of serial correlation in the standardized residuals. The panel also shows that for GM and IBM there is significant serial correlation in the squared standardized residuals.

2.3 Modelling the serial dependency in Realized Volatility - the IG-SV(p,q) model

From the results above it is clear that the simple specification corresponding to constant volatility is an insufficient description of the RV process, and it is clear that any type of framework we wish to apply for these series should be able to take account of this. In the present paper we assume that the underlying

¹The densities have been truncated at $RV = 5$ in order to better show the fit.

distribution is the IG and that the dynamics for RV_t can be described as

$$RV_t \sim IG(\sigma_t^2, a), \text{ where} \quad (7)$$

$$\sigma_t^2 = \omega + \sum_{i=1}^p \alpha_i RV_{t-i} + \sum_{j=1}^q \beta_j \sigma_{t-j}^2. \quad (8)$$

With this specification, the standardized RV's, $\overline{RV}_t = \frac{RV_t}{\sigma_t^2}$, are $IG(1, a)$ distributed and in the following we will refer to this model as the IG-SV(p,q) model. The specification in (7) – (8) can be estimated using a Maximum Likelihood approach. Note that in order to ensure that the variance remains positive we need to restrict the parameter space appropriately. Furthermore, using variance targeting we may fix $\omega = Var(R_t) * \left(1 - \sum_{i=1}^p \alpha_i - \sum_{j=1}^q \beta_j\right)$, where R_t is the asset returns, thus ensuring that the implied first moment of the RV series is matched to the actual level of the return variance.

While the model in (7) – (8) specifies what looks like an ARMA(p,q) process for the conditional level of the RV_t series it in fact corresponds more in style to a GARCH type process with RV_t playing the role of the innovations. In particular the same restrictions that are known from the GARCH framework apply in this model and will ensure positivity of the variance process as well as stationarity.

2.3.1 Estimation results

In Table 1 estimation results are reported for two models of RV. As mentioned above, Panel A presents the results for the simplest possible model with constant RV which corresponds to the special case in (8) where $p = q = 0$. In Panel B both p and q are allowed to differ from zero which means that the mean of the RV series becomes time varying. We report results here on the IG-SV(2,1) specification. However, in addition to this IG-SV(1,1) and IG-SV(1,2) models were also estimated but the preferred model is in all cases a IG-SV(2,1) model (the results are available from the author on request).

The results in Panel B from the IG-SV(2,1) specification indicates that with this more general specification significantly better results are obtained for all series. In particular, the panel shows that when the test statistics for misspecification are considered none of these are significant. Furthermore, all the parameters in the IG-SV(2,1) model are estimated significantly different from zero. In the right hand panels of Figure 3 the residual plots are shown for this model and it is clear from these that the fit by the model is very good.

Thus, we have shown that the RV series may be modelled well with the IG-SV(p,q) model. In particular the IG-SV(2,1) specification can accommodate the significant serial dependency observed in the RV data, and with the Inverse Gaussian distribution the model provides a very good fit to the data.

3 Returns

The previous section discussed ways to model the RV estimate of the variance process of financial assets. However, in order to be able to e.g. price derivative contracts a more elaborate model is needed. In particular it is necessary to specify the mean dynamics for the underlying asset in addition to the process for RV. In this section we start out by examining the empirical properties of the corresponding return series which we denote R_t . We also examine in detail the standardized returns, that is $R_t/\sqrt{RV_t}$. This analysis shows that the standardized returns have some particularly nice features. The observations we make enables us to propose a joint model for RV and asset return series.

3.1 Empirical regularities for returns and standardized returns

In the left hand panels of Figure 4 the time series of returns, R_t , are plotted for each of the assets considered. From these plots it can be noted that the series exhibit the same type of volatility clustering which is found in various other financial data series, and this finding confirms what is observed for the $\sqrt{RV_t}$ series in Figure 1. In addition to the volatility clustering, the right hand plots in the figure, which show the log density of the series with the Gaussian density superimposed, clearly show that the raw returns are far from being Gaussian distributed. In particular, from the plots it is clear that the unconditional distributions of the raw returns have significantly fatter tails than does the Gaussian distribution. Again this is a finding which has been well documented in the literature for various other financial data series.

However, a striking feature of the data is that when the returns are standardized with RV, the fit of the Gaussian distribution is much improved. This can be seen from the right hand plots of Figure 5 which show the log density plots of the standardized returns, $R_t/\sqrt{RV_t}$. Although departures from normality are still found, these are much less pronounced than for the raw series. The left hand plots in Figure 5 also confirms that there appears to be significantly less problems with volatility clustering once returns are standardized.

3.2 Distributional assumptions for the returns

It seems appropriate to base the joint modelling of return and RV processes on the above empirical regularities. Note that this is possible only because we have available the RV series from the high frequency intraday data. To be specific, we note that a nice result from the above findings is that we may maintain the distributional assumptions for the variance process in (7) – (8) and model the returns as being Gaussian distributed conditional on the RV process. Thus, a simply joint model for return and RV can be specified with a mean process and in this paper we choose to use the following specification

$$R_t = \mu_t + \sqrt{RV_t}\varepsilon_t, \text{ where } \varepsilon_t \sim N(0, 1). \tag{9}$$

We note that under the assumed dynamics the parameters of (7) – (8) can be estimated independently of (9). Likewise, given the return and RV series any parameters in (9) may be estimated easily from a simple linear regression. In particular it may be observed that, while the distribution of the returns, R_t , depends on the particular choice of the distribution for the variance, the standardized returns are Gaussian.

3.2.1 Empirical evidence for no price of Realized Volatility risk

While (9) allows for a very flexible specifications of μ_t an important special case would be to set

$$\mu_t = r + \lambda\sqrt{RV_t} - \frac{1}{2}RV_t. \quad (10)$$

In particular, if returns are Gaussian with variance equal to RV_t in (10) the parameter λ may be interpreted as the market price of RV risk. This specification corresponds closely to what has been used in the literature on relationships between risk and return. In our case the RV series are observable and therefore the specification allows us to test directly if the market requires a premium on risk or not. To do this we substitute (10) into (9) and rewrite it as follows

$$\begin{aligned} R_t &= r + \lambda\sqrt{RV_t} - \frac{1}{2}RV_t + \sqrt{RV_t}\varepsilon_t \Leftrightarrow \\ \frac{R_t - r - \frac{1}{2}RV_t}{\sqrt{RV_t}} &= \lambda + \varepsilon_t, \end{aligned} \quad (11)$$

where $\varepsilon_t \sim N(0, 1)$.

Hence, the market price of RV risk may be estimated as the constant term in(11). For the three stock returns under consideration the estimates, with P-values in parentheses, are 0.0238 (0.294) for GM, 0.0112 (0.611) for IBM, and 0.0531 (0.019) for MRK respectively. Thus we conclude that for the individual stock data used here there is no statistical evidence for a market price of RV risk.² While this is only one of many possible specifications we note that using other specifications in (10) like $\mu_t = r + \lambda RV_t - \frac{1}{2}RV_t$ or $\mu_t = r + \lambda - \frac{1}{2}RV_t$ lead to the same conclusions.

4 Option pricing using Realized Volatility

In the IG-SV(p,q) model the RV constitutes an additional factor of risk to investors and in general investors require a premium for being exposed to such risk. In particular, this is well known from the literature on stochastic volatility models where it is seen that this then changes the risk neutral dynamics to be used for e.g. pricing purposes. However, as we have seen above it does in fact not appear that RV risk is priced.

²We performed the same regression for the Standard & Poor's 100 Index, and in this case a highly significant coefficient was obtained. Thus, we refrain from applying our model to indices.

In the following we make this an explicit assumption of the modelling framework which, together with the stated dynamics, allows us to formally derive the risk neutralized dynamics to be used for option pricing. We then explain how the option pricing model may be implemented. We end this section by analyzing the properties of the IG-SV(p,q) option pricing model in comparison to the classical Black-Scholes-Merton framework which can be obtained as a special case of our proposed framework.

4.1 The IG-SV(p,q) option pricing model

In order to formally derive the IG-SV(p,q) option pricing model we make explicit the following two assumptions which build on the empirical findings reported in Sections 2 and 3 above:

Assumption 1 Under the physical measure asset returns, R_t , are Gaussian conditional on Realized Volatility, RV_t , which in turn follows the IG-SV(p,q) model. Formally we have the following dynamics:

$$R_t = \mu_t + \sqrt{RV_t}\varepsilon_t, \text{ where } \varepsilon_t \sim N(0, 1), \text{ and where} \quad (12)$$

$$RV_t \sim IG(\sigma_t^2, a), \text{ with} \quad (13)$$

$$\sigma_t^2 = \omega + \sum_{i=1}^p \alpha_i RV_{t-i} + \sum_{j=1}^q \beta_j \sigma_{t-j}^2. \quad (14)$$

Assumption 2 Realized Volatility is observable but is not a traded asset. Furthermore, investors require no premium for being exposed to Realized Volatility risk when holding financial assets.

Assumption 1 simply specifies formally the joint distribution of returns and RV. As we have seen in Section 2.3.1 the specification in (13) – (14) appears to be reasonable for the asset returns considered in this paper. Furthermore, in Section 3 we argued in favor of the conditional distribution in (12). With respect to Assumption 2 the first part simply states that investors have access to the RV measure although they cannot trade it. We believe this to be a quite realistic assumption. The second part, although it may appear controversial, is supported by the data on individual stocks used in this paper as we have seen in Section 3.2.1.

When it comes to option pricing Assumption 2 is very convenient since it implies that the risk neutral RV dynamics are the same as the physical dynamics. Thus, we should use the dynamics specified in (13) – (14) for option pricing purposes. However, since the underlying asset is a traded asset its return dynamics needs to be changed under the risk neutral dynamics such that the expected return equals the risk free interest rate. For this to be the case we must have

$$\exp(r) = E^Q \left[\exp \left(\mu_t + \sqrt{RV_t}\varepsilon_t \right) \middle| \mathcal{F}_{t-1} \right]. \quad (15)$$

Assuming that μ_t is measurable with respect to the time $t - 1$ information set we may rewrite this to yield

$$\mu_t = r - \ln E^Q \left[\exp \left(\sqrt{RV_t}\varepsilon_t \right) \middle| \mathcal{F}_{t-1} \right], \quad (16)$$

which is the mean specification to be used in the risk neutralized dynamics.

We are now ready to formally characterize the risk neutral dynamics to be used for option pricing purposes in the IG-SV(p,q) model. These are given by the following proposition:

Proposition 1 *Under Assumption 1 and 2 the risk neutral dynamics for Realized Volatility remains within the IG-SV(p,q) family and the process is characterized by (13) – (14). However, under the risk neutral dynamics the appropriate mean specification for the asset return is given by*

$$R_t = r - \ln E^Q \left[\exp \left(\sqrt{RV_t} \varepsilon_t \right) | \mathcal{F}_{t-1} \right] + \sqrt{RV_t} \varepsilon_t, \quad (17)$$

where $\varepsilon_t \sim N(0, 1)$.

4.1.1 Implementing the IG-SV(p,q) option pricing model

In order to be able to implement the IG-SV option pricing model it is necessary to provide methods for calculating the risk neutralized mean equation in (17) in Proposition 1. We first consider the special case with deterministic RV_t , e.g. when $\alpha_i = 0$ for all i , of which the constant volatility is a special case. In this situation it follows that conditional on time $t-1$ information $\sqrt{RV_t} \varepsilon_t$ is Gaussian distributed with mean zero and variance RV_t under the risk neutralized dynamics. From the properties of the log normal distribution it then follows that $E^Q \left[\exp \left(\sqrt{RV_t} \varepsilon_t \right) | \mathcal{F}_{t-1} \right] = \exp \left(\frac{1}{2} RV_t \right)$, and in this case the risk neutralized mean equation in (17) is given by

$$R_t = r - \frac{1}{2} RV_t + \sqrt{RV_t} \varepsilon_t. \quad (18)$$

In the more general case it may appear that this calculation is overly complicated and that it requires some sort of approximation. However, this is not the case since the evaluation of $E^Q \left[\exp \left(\sqrt{RV_t} \varepsilon_t \right) | \mathcal{F}_{t-1} \right]$ can be performed using the appropriate Moment Generating Function, or MGF. To be specific, recall that the MGF of a random variable X is specified as

$$M_X(t) = E \left[\exp(tX) \right]. \quad (19)$$

Thus, assuming the MGF exist and can be evaluated at $t = 1$ it can be used to evaluate the expectation to be used for risk neutralization since $E \left[\exp(X) \right] = M_X(1)$. In fact, what is needed is the log of this expectation which corresponds to the cumulant generating function.

In the IG-SV(p,q) model we recall that $RV_t \sim IG(\sigma_t^2, a)$ and $\varepsilon_t \sim N(0, 1)$. Hence, if we define $X_t = \sqrt{RV_t} \varepsilon_t$ by construction it follows that X_t is a $NIG(0, \delta, a, 0)$ distributed variable with mean zero and variance σ_t^2 . Moreover, since the variance of a $NIG(0, \delta, a, 0)$ distributed variable, X , is given by

$$Var(X) = \frac{\delta^2}{a}, \quad (20)$$

it follows that $\delta = \sqrt{\sigma_t^2 a}$. From e.g. Jensen & Lunde (2001) we know that the MGF of a $NIG(\mu, \delta, a, b)$ distributed variable specified in terms of the scale invariant parameters is given by

$$M_{X \sim NIG(\mu, \delta, a, b)}(t) = \exp \left\{ \mu t + a \left(\sqrt{1 - (b/a)^2} - \sqrt{1 - \left(\frac{b + \delta t}{a} \right)^2} \right) \right\}. \quad (21)$$

In the present setting $\mu = 0$, $\delta = \sqrt{\sigma_t^2 a}$, and $b = 0$, and for X_t we therefore have that

$$M_{X_t \sim NIG(0, \sqrt{\sigma_t^2 a}, a, 0)}(1) = E^Q [\exp(X_t) | \mathcal{F}_{t-1}] = \exp \left\{ a \left(1 - \sqrt{1 - \sigma_t^2/a} \right) \right\}. \quad (22)$$

We note that as a tends to infinity the limit of (22) is $\exp(\frac{1}{2}\sigma_t^2)$ which corresponds to the correction in the Gaussian case.

Thus, in the general case the mean equation in the risk neutralized dynamics of the IG-SV(p,q) option pricing model has a very simple form which we state formally in the following corollary:

Corollary 1 *Under Assumptions 1 and 2 the logarithm of the expectation $E^Q [\exp(\sqrt{RV_t} \varepsilon_t) | \mathcal{F}_{t-1}]$ needed to ensure appropriate risk neutralization in Proposition 1 can be calculated explicitly using the following relationship when this is well defined*

$$\ln E^Q \left[\exp(\sqrt{RV_t} \varepsilon_t) | \mathcal{F}_{t-1} \right] = a \left(1 - \sqrt{1 - \sigma_t^2/a} \right). \quad (23)$$

In particular, the mean equation in (17) in Proposition 1 is given by

$$R_t = r - a \left(1 - \sqrt{1 - \sigma_t^2/a} \right) + \sqrt{RV_t} \varepsilon_t. \quad (24)$$

4.2 Properties of the Realized Volatility Option Pricing Model

Before we put the IG-SV(p,q) option pricing model to the test empirically it is of interest to examine the properties of the model in more detail. In particular, it is useful to compare the model to existing alternatives, like e.g. the constant volatility Gaussian model. Doing so will allow us to assess whether or not the suggested model will eventually be able to explain such alternative model shortcomings. However, because of the lack of closed form solutions to the option pricing problem this can not be done analytically. Instead we have to analyze the performance of the RV option pricing model using numerical procedures.

In this paper we suggest to do this using Monte Carlo simulation. In particular, we use the derived risk neutral dynamics to simulate a large number of future paths of the price of the underlying and of the RV and use these to price a set of artificial options with varying strike prices, K , and maturities, T . Option valuation using Monte Carlo simulation is a well established procedure and has been used at least since Boyle (1977) for European style options. However in the present setting we need to take account of the early exercise feature in the American style options. We do this using the Least Squares Monte Carlo (LSM) method of Longstaff & Schwartz (2001) where the early exercise value is estimated using the cross sectional information available in the simulation.

4.2.1 Realized Volatility option pricing and the Black-Scholes-Merton benchmark

In Table 2 we report price estimates for the Gaussian-CV model, the IG-CV model, and the IG-SV(2,1) model, where the first model corresponds to the Black-Scholes-Merton benchmark. The columns headed “Rel Diff” reports the mispricing which would be observed with the Gaussian benchmark model if the dynamics of the true world were IG-CV or IG-SV(2,1), respectively. The parameters used in this study correspond roughly to the average estimates in Table 1. In particular, we chose $a = 1.4$ for the underlying distribution and $\alpha_1 = 0.21$, $\alpha_2 = -0.16$, and $\beta = 0.93$ in the IG-SV(1,2) specification. In all cases we fix ω to yield an annualized level of volatility equal to 25%. For the simulation we use 20,000 paths and in the cross sectional regressions we use powers and cross products of asset prices and RV of order less than or equal to two. The reported prices and standard errors are calculated from 100 independently generated simulations.

Comparing first the two constant volatility versions to each other the table clearly shows that if the true dynamics were in fact of the IG type the classical Gaussian Black-Scholes-Merton model would severely underprice short term out of the money options. In fact this is what is often found empirically and it has become known as the “smile” effect in option prices. Thus, this finding indicates the potential of the IG framework to improve on existing models shortcomings. However, the table also shows that this feature of the IG-CV model vanishes quite quickly as the time to maturity is increased and for options with three or more months to maturity it is negligible.

For the IG-SV(2,1) model the overall conclusions are the same. However the table also shows that if the true dynamics for RV include time variation in the mean in addition to being IG distributed the underpricing by the classical Gaussian model would be even more pronounced and this is particularly so for the short term call options. Furthermore, the mispricing is present for longer maturities and for the put options the underpricing remains approximately 26% for the out of the money options even when time to maturity is the equivalent of one month. Note though that as time to maturity increases the pricing difference vanishes as it was the case with the IG-CV formulation although the effect remains present even with three months option. We note that it is to be expected that the effects vanish for long term options as a central limit theorem starts to apply for the combined sum of innovations which approaches a Gaussian variable.

5 Interdaily “limit” models

The IG-SV(p,q) model we have suggested assumes that RV is an observable variable which can be modelled as being IG distributed with a dynamic specification for the conditional mean. The joint model of RV and returns is based on the observation that when returns are standardized with RV they are approximately Gaussian. Hence in this framework asset returns are NIG distributed with a time varying volatility. Thus,

when analyzing the incremental value of the high frequency data a valid benchmark would be a framework which shares these characteristics but which is based on the subset of information contained in interdaily asset returns only.

We suggest that an obvious benchmark is the GARCH framework. In particular, this is a discrete time framework in which asset returns are allowed to have time varying volatility. Furthermore, although the classical GARCH model assumes a Gaussian conditional distribution this may be relaxed and instead the NIG distribution can be used. In this section we therefore compare models within this framework to IG-SV models both in terms of estimation results and option pricing characteristics. We end the section by reinterpreting the discrete time GARCH framework as a multiplicative Errors in Variables model for the IG-SV framework when RV is unobservable and the squared returns are used as a proxy.

5.1 NIG GARCH models as an appropriate benchmarks

One particular set of models which is quite closely related to the IG-SV(p,q) model above but is based only on interdaily return data is the family of NIG GARCH models developed in Stentoft (2006). In the simplest possible formulation within this framework the return dynamics are assumed to be

$$R_t = \mu_t + \sqrt{h_t}\varepsilon_t, \text{ where } \varepsilon_t \sim NIG(a, b), \quad (25)$$

and where the volatility, h_t , follows a GARCH(p,q) type process such as

$$h_t = \omega + \sum_{i=1}^q \alpha_i h_{t-i} \varepsilon_{t-i}^2 + \sum_{j=1}^p \beta_j h_{t-j}. \quad (26)$$

In (25) the notation $NIG(a, b)$ denotes the standardized NIG distribution with mean zero and unit variance where a governs the leptokurtosis and b governs the skewness (see also Stentoft (2006)).

The reason for the similarity is first of all that in the NIG GARCH model the innovations, $\sqrt{h_t}\varepsilon_t$, are NIG distributed. However, we recall that this is the same type of distribution which the innovations will have in the IG-SV(p,q) model according to (9) under the above assumptions on the dynamics for the RV. Secondly, we note that when comparing the dynamic specifications in (8) and (26) in the general case the expressions for model volatility are potentially quite similar. In particular, in general the NIG GARCH model has the potential to generate time varying volatility as it is the case for the IG-SV model.

Thus, although the two models are equivalent only when $p = q = 0$ in both (8) and (26) in which case both specification leads to a constant volatility model with NIG innovations, we may conclude that within either of the two frameworks models can be specified with the same overall return characteristics. Furthermore, since the NIG GARCH framework is based on the interdaily return data these models rely only on a subset of the information set used in the IG-SV framework. Therefore this type of models serve as a natural set

of benchmark models when the goal is to analyze the incremental value of using RV data for e.g. option pricing.

One might argue that since the IG-SV model may be interpreted as a stochastic volatility model under additional assumptions on the RV series it should be compared to such models and not to GARCH type models which only rely on one stochastic factor. However, in general interdaily stochastic volatility models treats volatility as an unobservable state variable which is not the case in the IG-SV model where RV is assumed observable. Hence the stochastic volatility models rely on fundamentally different assumptions and would require a change in the applied estimation technique. For this reason we leave this comparison for future research.

5.1.1 Estimation results for the NIG GARCH model relative to the IG-SV model

In Table 3 results are provided for two models from the NIG GARCH(p,q) framework estimated on daily return data for the same period as was used above for the IG-SV model. To be specific, we report results for the simple constant volatility specification as well as for a model with time varying volatility with $p = 1$ and $q = 2$, that is with the same number of parameters as the chosen specifications for the models based on RV. We report only the results with $b = 0$ as this parameter was insignificantly different from zero for all asset returns series for this period.

In the estimation we set μ_t in (25) equal to

$$\mu_t = r - \ln E \left[\exp \left(\sqrt{h_t} \varepsilon_t \right) | \mathcal{F}_{t-1} \right]. \quad (27)$$

In so doing we ensure that the expected asset return equals the risk free interest rate. We note that this corresponds to what is referred to as the feasible specification of the NIG GARCH model in Stentoft (2006). We note that (22) may be used to show that the expectation in (27) can be calculated explicitly from the MGF of a NIG distributed variable as it was the case for the IG-SV(p,q) model with h_t replacing σ_t^2 . This differs from the approximations outlined in ?? but is similar to what is suggested in Christoffersen, Elkamhi, Feunou & Jacobs (2008).

The results in Panel A of the table shows that a constant volatility model is unable to accommodate the properties of the data. In particular, the panel shows that for all the time series significant correlation in the squared residuals are found which together with the significant statistics for ARCH1-5 indicates that there are time variation of e.g. the GARCH type in the data series. On the other hand, Panel B shows that a NIG GARCH(1,2) model is able to accommodate the time varying volatility, and for this model none of the tests for misspecification are significant.

When comparing the estimation results for the NIG GARCH model to those obtained with the IG-SV model we observe that there are definite similarities. In particular, for the two models the a parameter

is estimated at very low values indicating a high degree of leptokurtosis in the return data. Furthermore, although there are differences in the point estimates these is not significant though in part because of the large standard errors on the estimates from the NIG GARCH models. With respect to the parameters of the volatility specification the two models are also similar in terms of the implied persistence in the variance. However, the actual point estimates do differ and this is so especially for the α parameters.

5.2 Option pricing with NIG GARCH models

In the NIG GARCH model in (25) and (26) the dynamics to be used for option pricing may be derived from Duan (1999) as it was done in Stentoft (2006). To be specific, the model we propose is a special case of the general NIG GARCH framework used in Stentoft (2006) with $\lambda = 0$ and with a symmetric distribution since $b = 0$. Thus, it follows that the dynamics of the risk neutralized process are

$$R_t = r - \ln E^Q \left[\exp \left(\sqrt{h_t} \varepsilon_t \right) | \mathcal{F}_{t-1} \right] + \sqrt{h_t} \varepsilon_t, \text{ where } \varepsilon_t \sim NIG(a, 0) \text{ and} \quad (28)$$

$$h_t = \omega + \sum_{i=1}^q \alpha_i h_{t-i} \varepsilon_{t-i}^2 + \sum_{j=1}^p \beta_j h_{t-j}. \quad (29)$$

Again we note that (22) may be used to show that the expectation in (28) can be calculated explicitly from the MGF of a NIG distributed variable as it was the case for the IG-SV(p,q) model with h_t replacing σ_t^2 .

As it is the case with the model under the physical measure we observe that the model in (28) – (29) looks similar to the IG-SV option pricing model and that they are equivalent when RV_t and h_t are constant. In the general case however the two models differ although they share the same overall properties under the risk neutral measure. In particular, the innovations under the risk neutral dynamics are in both cases NIG distributed and both models will generate time varying volatility.

Once again it could be argued that another valid benchmark would be a stochastic volatility model. In particular, numerous continuous time stochastic volatility option pricing models have been developed starting with the classical contributions of Hull & White (1987), Johnson & Shanno (1987), Scott (1987), and Wiggins (1987). One obvious advantage with these models is that the assumption of continuous trading allows for elegant derivation of the risk neutral dynamics. However, when actually applied most of these models again treat volatility as an unobservable state variable which complicates significantly the prediction step needed for option pricing. In particular, this may require e.g. the full reprojection machinery associated with the EMM procedure (see Gallant & Tauchen (1998)). For this reason we again leave this comparison for future research.

5.2.1 Option pricing results with the NIG GARCH model relative to the IG-SV model

In Table 4 we report results corresponding to those in Table 2 but using the NIG GARCH model. Thus, comparing the two tables enables us to gauge the potential differences between models based on the RV series derived from high frequency data and models relying only on the daily return series. For the NIG GARCH model the same a is used for the distribution and the volatility parameters used are $\alpha_1 = 0.08$, $\alpha_2 = -0.06$, and $\beta = 0.96$. These parameters corresponds to what is found when estimation the models and imply the same persistence of 0.98 in the volatility dynamics as was the case with the IG-SV model.

When analyzing the pricing performance the table shows that the same overall conclusions can be drawn as with the IG-SV model. In particular, if the true dynamics were in fact those from a NIG GARCH specification a Gaussian constant volatility model would severely underprice particularly short term out of the money options. Furthermore, although this effect vanishes quickly with a constant volatility NIG model when GARCH effects are present the mispricing persists for longer maturities.

When comparing between Tables 2 and 4 we note that overall the pricing performances are similar for the two constant volatility specifications and for the models with time varying volatility. However, this is in fact to be expected since the relevant models share the same general properties, and as such this validates the use of the NIG GARCH framework as a benchmark for comparison. In particular, the pricing results suggests that any differences observed in pricing performance will be caused by different informational content in the series in general and different parameter estimates in particular rather than by differences in the actual pricing methodology.

5.3 The NIG GARCH model as a multiplicative Errors in Variables model

Until now we have argued in favour of using the NIG GARCH framework as a benchmark because such models are able to generate the same characteristics as models within the IG-SV framework while relying on a restricted information set. We now formally justify this by interpreting the NIG GARCH framework as a multiplicative Errors in Variables model for the IG-SV model which uses noisy observations of RV.³ To be specific, consider the simple case when $\mu_t = 0$ in the mean equation in (25) of the NIG GARCH model. Thus, in this model the innovation term, $h_t \varepsilon_t^2$, which is used in (26), corresponds to the squared interdaily return. While this may be interpreted as an estimate of daily realized volatility, it is potentially a very noisy estimate.

From the literature on Errors in Variables models we know that the parameter estimates associated with the affected variables may be severely biased towards zero. In the following we examine this in the case of the proposed models. In particular, in the next section we perform a Monte Carlo study which confirms

³We thank Éric Jacquier for suggesting this interpretation.

that such a bias is found if the NIG GARCH models are estimated on series which are generated by IG-SV models. We end this section by providing estimation results on NIG GARCH models which are augmented to take into consideration the RV data.

5.3.1 A Monte Carlo study of the Errors in Variables problem for the NIG GARCH model

In order to examine the anticipated effects on the parameter estimates we perform a small Monte Carlo study where we estimate the NIG GARCH model on data series generated from the IG-SV model using both the “Errors in Variable” specification and a specification which augments the conditional variance equation of the GARCH model with the actual RV series. In particular, the latter specification uses

$$h_t = \omega + \sum_{i=1}^{q_{RV}} \alpha_{RV,i} RV_{t-i} + \sum_{i=1}^{q_{EiV}} \alpha_{EiV,i} h_{t-i} \varepsilon_{t-i}^2 + \sum_{j=1}^p \beta_j h_{t-j}, \quad (30)$$

for the conditional variance equation instead of (25).

For the simulation we use the average values of the parameters from estimation of the preferred IG-SV(2,1) specification. In particular, we set $\alpha_1 = 0.21$, $\alpha_2 = -0.16$, $\beta_1 = 0.93$, and fix ω to yield an annualized level of volatility equal to 25%. Furthermore, we choose $a = 1.4$ which again corresponds to the empirical average for the three series. A sample of 10,000 observations is used and this is repeated 100 times. For the estimation we choose an equivalently parameterized NIG GARCH(1,2) specification which is estimated on the last 5,000 observations from the simulated returns and RV’s discarding the first 5,000 observations.

The results of the Monte Carlo study are shown in Table 5. Under the heading “Squared Returns” the estimation result using the normal GARCH specification are shown and we observe that while the parameter estimates have the correct sign and are significantly different from zero they are significantly biased towards zero as we would expect. This is not so if the actual RV series is used as the results under the heading “Realized Volatility” show. In particular, we observe that when the actual RV series is used in the GARCH framework the parameter estimates are insignificantly different from those of the IG-SV model. The last set of results shows that when using squared returns in addition to the RV series the squared returns are insignificant and the results confirm that with the augmented NIG GARCH framework the correct dynamics are in fact retrieved. Again we interpret these results as a validation of our choice of the GARCH frameworks as a relevant benchmark model.

5.3.2 Estimation results for the NIG GARCH model augmented with the Realized Volatility

Based on the findings above in Table 6 we report estimation results for the NIG GARCH(1,2) model when the specification is augmented with the actual observed RV series. That is, we specify the conditional volatility as in (30) with $q_{RV} = q_{EiV} = 2$ and $p = 1$.

The first observation we make is that the estimates relating to the squared return parameters are insignificant. This holds for all stocks and shows that the explanatory power of squared returns in terms of the GARCH process is completely subsumed by the RV series. Furthermore, the results may be compared directly with those in Panel B of Table 3. When doing this we observe that the likelihood values are increased significantly when using the RV as explanatory variables in the volatility specification, which indicates that this model is more appropriate than the original GARCH specification. This conclusion is supported by comparing the values of the Schwarz Information Criteria which for all stocks is minimized with the augmented model.

Finally, the estimation results for the augmented NIG GARCH model may also be compared to those for the IG-SV model in Table 1. When doing so we find that the actual estimates are quite similar. In particular, comparing the tables we observe that the estimates of $\alpha_{RV,1}$ and $\alpha_{RV,2}$ are in fact insignificantly different from the corresponding estimates in the IG-SV model which was not the case when only the squared returns were used. However, the table also shows that the parameters are estimated with more uncertainty when the GARCH specification is used instead of the IG-SV specification. This indicates that the latter method is in fact more efficient and validates our modelling approach within the IG-SV framework.⁴

6 Empirical Performance of the Realized Volatility Option Pricing Model

In this section we provide results on the empirical performance of the RV option pricing model. To be specific we will use the model to price options on the three individual stocks considered previously for the year 1995. The option data we use consists of an end of day observation for all contracts which were traded more than five times during the day. The parameters needed for the simulation are reestimated daily using only historically available information and in the pricing algorithm we use the method outlined in Section 4.2. We use the current LIBOR rate as the risk free interest rate and assume that future dividend payments are known in advance and that they spill over fully on the simulated future stock prices.

In the next section we provide results for the performance of the various RV-SV(p,q) models which shows that the time varying specification is important. In the second section the results are benchmarked against a equivalently parameterized specification from the NIG GARCH framework. This section shows that the RV based option pricing model outperforms the models which are based only on interdaily data and hence provides evidence on the value of using high frequency data for option pricing purposes.

⁴Furthermore, the augmented GARCH framework does not include dynamics for RV and hence it would be impossible to simulate from this process when performing e.g. option pricing.

6.1 The empirical performance of the Realized Volatility Option Pricing Model

In Table 7 we report the absolute pricing errors for different IG-SV(p,q) option pricing models using two standard measures from the literature. In particular, letting P_k and \tilde{P}_k denote the k 'th observed price and the k 'th model price these are the relative mean bias, $RBIAS \equiv K^{-1} \sum_{k=1}^K \frac{(\tilde{P}_k - P_k)}{P_k}$, and the relative mean squared error, $RSE \equiv K^{-1} \sum_{k=1}^K \frac{(\tilde{P}_k - P_k)^2}{P_k^2}$. The reason for using these relative metrics is that the actual option prices varies significantly between the in the money and out of the money options. Thus, if absolute metrics are used much more weight would be put on pricing the in the money options correctly than on pricing the out of the money options.

The first thing we note from the table is that overall the constant volatility model has the largest pricing errors for all underlying assets, and this holds when the put and call options are considered separately. On the other hand the models with stochastic RV dynamics, like the IG-SV(1,1) or IG-SV(2,1) specifications, perform much better. In particular the IG-SV(2,1) specification is the best performing model for both metrics in all but one situation the exception being when using the RBIAS metric for IBM call options. Thus, in terms of the overall pricing performance based on RV we find strong evidence in favour of the suggested model with a time varying specification.

6.1.1 Performance across moneyness

In Table 8 we report the overall pricing performance of the various IG-SV(p,q) models across moneyness. We define moneyness as $M = S / (K * \exp(-r * T))$ where K is the strike price and T is the time to maturity. We let the cut off be $\pm 2\%$ and thus define out of the money call options as those options with $M < 98\%$ whereas the in the money call option has $M > 102\%$. For the put options this is the opposite.

The table clearly confirms that large improvements are found with time varying models over constant volatility models consistently across moneyness. However, the table also indicates that the improvements are particularly large for out of the money options. This is in line with what we would expect based on the study of the properties of the IG-SV(p,q) option pricing model in Section 4.2. The improvements are particularly large for IBM where the pricing errors from the IG-SV(2,1) model are only a fraction of those from the constant volatility benchmark when using the RBIAS metric for out of the money options.

6.2 The performance when benchmarked against purely return driven models

In Table 9 we report the relative pricing errors for the IG-SV option pricing model when compared to the NIG GARCH option pricing model with the same number of parameters. To be specific, the table compares the IG-SV(0,0) model to the NIG GARCH(0,0) model, the IG-SV(1,0) model to the NIG GARCH(0,1) model, the IG-SV(1,1) model to the NIG GARCH(1,1) model, and the IG-SV(2,1) model to the NIG GARCH(1,2)

model. In all panels the ratio expressed is the IG-SV pricing error divided by the corresponding NIG GARCH pricing error minus one. Thus, negative cell numbers indicate that the SV model is the best performing model for that particular combination of option, metric, and number of parameters. On the other hand a positive value would indicate that the return based model is in fact the best performing one. The actual value of the cell corresponds to the improvement in performance respectively the deterioration in performance when using RV data.

The first thing that may be observed from the table is that in the majority of the combinations of metrics and types of options the IG-SV specifications outperforms the corresponding return based NIG GARCH specification. Although the difference is minor for the constant volatility models it increases with the complexity of the model, and when the IG-SV(2,1) models are compared to the corresponding NIG GARCH(1,2) models the improvements are all large. In particular, this is the case for IBM where the RBIAS is reduced with more than 90% when using the intraday RV data in the IG-SV model. The results holds when put and call options are considered individually although the improvements are slightly smaller for the call options than for the put options particular in terms of the RBIAS metric.

6.2.1 Performance across moneyness

In Table 10 we report the relative pricing performance of the various IG-SV(p,q) models across moneyness. The table clearly confirms that when benchmarked against models using only daily data large improvements are found with the IG-SV models across moneyness. However, the table also indicates that the improvements are relatively more spectacular when pricing options which are either out of the money or in the money. In terms of option pricing this would indicate that the IG-SV models generate a more pronounced smile effect than the corresponding GARCH models which confirms what was found in the Monte Carlo study in Section 4.2.

7 Conclusion

In this paper we propose to model Realized Volatility using the Inverse Gaussian distribution as the conditional distribution and with a time varying mean specification. We denote this model the IG-SV(p,q) model. We then derive the appropriate model to be used for option pricing and we analyze the model's properties through a Monte Carlo study indicating that it may explain some of the mispricings found when using traditional option pricing models based on interdaily data. The model we have proposed jointly for Realized Volatility and returns implies that returns are Normal Inverse Gaussian and we argue that a natural benchmark for the IG-SV model are models within the NIG GARCH framework. We show that these models may be interpreted as a multiplicative Errors in Variables model for the IG-SV(p,q) model. We then

provide an empirical examination of the pricing performance of the IG-SV model and show that it performs significantly better than models within the NIG GARCH framework which only rely on interdaily return data. The results provide evidence of the value of using high frequency data for option pricing purposes.

While the results of this paper are very promising for the use of Realized Volatility estimates for option pricing several issues may warrant further examination. In particular, the measure used here for Realized Volatility is the simplest possible available. Hence it would be interesting to analysis the effect of using other more robust measures for option pricing. Furthermore, the model used here made the simplifying assumption that Realized Volatility risk is not priced in the market. Although this assumption may be reasonable considering the return data it is possible that the price of Realized Volatility risk can be estimated better by using historical observations on actual options. Finally, the present paper only considers options on three underlying assets and this for a limited period. It is obviously of interest to expand the analysis in order to examine the robustness of the findings.

References

- Andersen, T. G., Bollerslev, T., Diebold, F. X. & Labys, P. (2001), ‘The Distribution of Exchange Rate Volatility’, *Journal of the American Statistical Association* **96**, 42–55.
- Andersen, T. G., Bollerslev, T., Diebold, F. X. & Labys, P. (2003), ‘Modeling and Forecasting Realized Volatility’, *Econometrica* **71**(2), 579–625.
- Barndorff-Nielsen, O., Hansen, P. R., Lunde, A. & Shephard, N. (2006), ‘Designing realised kernels to measure the ex-post variation of equity pricis in the presence of noise’, *Unpublished paper: Nuffield College, Oxford*.
- Barndorff-Nielsen, O. & Shephard, N. (2002), ‘Econometric analysis of realized volatility and its use in estimating stochastic volatility models’, *Journal of the Royal Statistical Society, Series B* **64**, 253–280.
- Black, F. & Scholes, M. (1973), ‘The Pricing of Options and Corporate Liabilities’, *Journal of Political Economy* **81**, 637–654.
- Bollerslev, T. (1986), ‘Generalized Autoregressive Conditional Heteroskedasticity’, *Journal of Econometrics* **31**, 307–327.
- Bollerslev, T., Chou, R. Y. & Kroner, K. F. (1992), ‘ARCH Modelling in Finance’, *Journal of Econometrics* **52**, 5–59.
- Boyle, P. P. (1977), ‘Options: A Monte Carlo Approach’, *Journal of Financial Economics* **4**, 323–338.

- Christoffersen, P., Elkamhi, R., Feunou, B. & Jacobs, K. (2008), ‘Option Valuation with Conditional Heteroskedasticity and Non-Normality’, *SSRN Working Paper* (<http://ssrn.com/abstract=961512>).
- Clark, P. K. (1973), ‘A subordinated stochastic process model with finite variance for speculative prices’, *Econometrica* **41**(1), 135–155.
- Duan, J.-C. (1999), ‘Conditionally Fat-Tailed Distributions and the Volatility Smile in Options’, *Mimeo, Hong Kong University of Science and Technology*.
- Engle, R. F. (1982), ‘Autoregressive Conditional Heteroscedasticity With Estimates of The Variance of United Kingdom Inflation’, *Econometrica* **50**(4), 987–1007.
- Forsberg, L. & Bollerslev, T. (2002), ‘Bridging the Gap Between the Distribution of Realized (ECU) Volatility and ARCH Modeling (of the EURO): The GARCH Normal Inverse Gaussian Model’, *Journal of Applied Econometrics* **17**, 535–548.
- Gallant, R. & Tauchen, G. (1998), ‘Reprojecting Partially Observed Systems With Application to Interest Rate Diffusions’, *Journal of the American Statistical Association* **93**(441), 10–24.
- Hansen, P. R. & Lunde, A. (2005a), ‘A Forecast Comparison of Volatility Models: Does Anything Beat a GARCH(1,1)’, *Journal of Applied Econometrics* **20**, 873–889.
- Hansen, P. R. & Lunde, A. (2005b), ‘A Realized Variance for the Whole Day Based on Intermittent High-Frequency Data’, *Journal of Financial Econometrics* **3**(4), 525–554.
- Heston, S. L. (1993), ‘A Closed-Form Solution for Options with Stochastic Volatility with Applications to Bond and Currency Options’, *Review of Financial Studies* **6**(2), 327–343.
- Hull, J. & White, A. (1987), ‘The Pricing of Options on Assets with Stochastic Volatilities’, *Journal of Finance* **42**(2), 281–300.
- Jensen, M. B. & Lunde, A. (2001), ‘The NIG-S&ARCH model: a fat-tailed, stochastic, and autoregressive conditional heteroskedastic volatility model’, *Econometrics Journal* **4**, 319–342.
- Johnson, H. & Shanno, D. (1987), ‘Option Pricing When the Variance is Changing’, *Journal of Financial and Quantitative Analysis* **22**(2), 143–151.
- Longstaff, F. A. & Schwartz, E. S. (2001), ‘Valuing American Options by Simulation: A Simple Least-Squares Approach’, *Review of Financial Studies* **14**, 113–147.
- Merton, R. C. (1973), ‘Theory of Rational Option Pricing’, *Bell Journal of Economics and Management Science* **4**, 141–183.

- Poon, S.-H. & Granger, C. W. J. (2003), 'Forecasting Volatility in Financial Markets: A Review', *Journal of Economic Literature* **XLI**, 478–539.
- Scott, L. O. (1987), 'Option Pricing When the Variance Changes Radomly: Theory, Estimation, and an Application', *Journal of Financial and Quantitative Analysis* **22**(4), 419–438.
- Stentoft, L. (2006), 'American Option Pricing using GARCH models and the Normal Inverse Gaussian distribution', *SSRN Working Paper* (<http://ssrn.com/abstract=946002>) .
- Wiggins, J. B. (1987), 'Option Values Under Stochastic Volatility: Theory and Empirical Estimates', *Journal of Financial Economics* **19**, 351–372.
- Zhang, L., Mykland, P. A. & Aït-Sahalia, Y. (2005), 'A Tale of Two Time Scales: Determining Integrated Volatility with Noisy High-Frequency Data', *Journal of the American Statistical Association* **100**, 1394–1411.

Tables and Figures

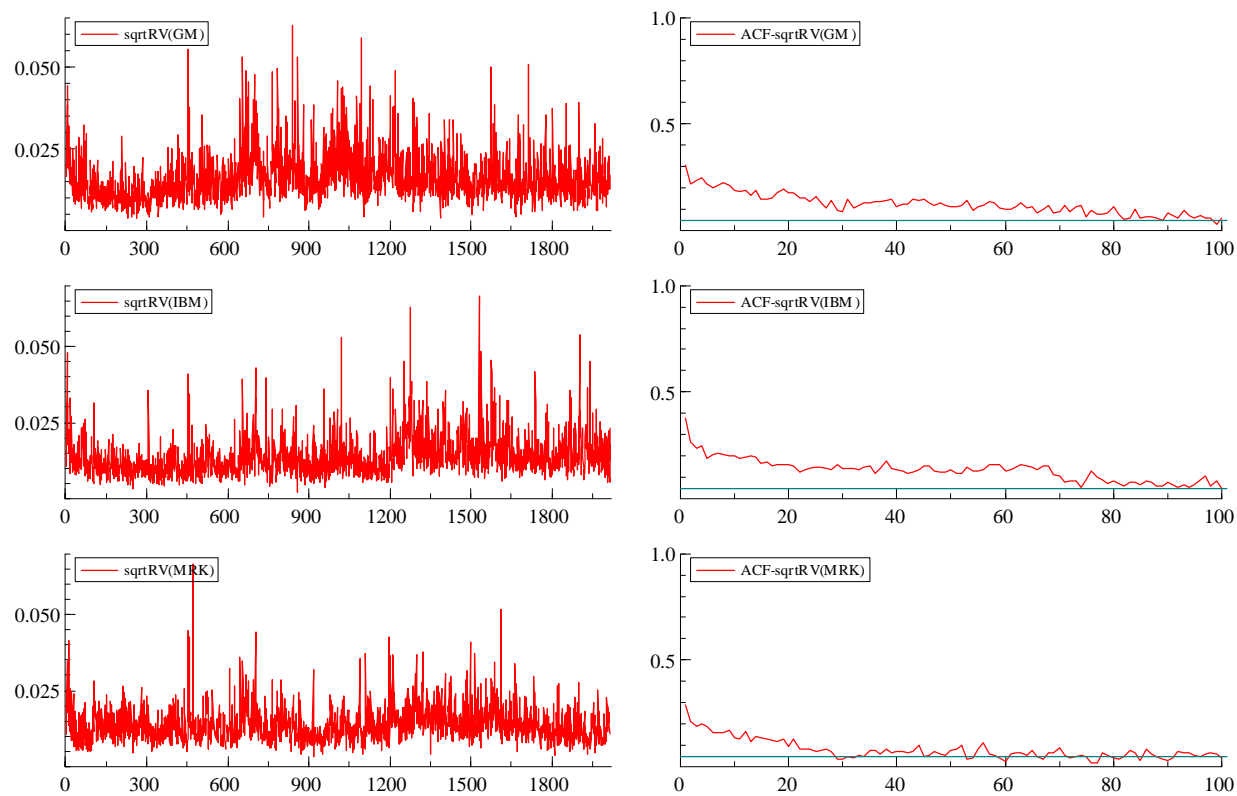


Figure 1: This figure shows time series plots of $\sqrt{RV_t}$ in the left hand plots and of the autocorrelations for the first 100 lags of $\sqrt{RV_t}$ in the right hand plots. From top to bottom the plots are for GM, for IBM, and for MRK.




Figure 2: This figure shows plots of the Inverse Gaussian density, $f_{IG}(z; 1, a)$ for three different value of a . The solid line corresponds to $a = 0.5$, the solid dotted line corresponds to $a = 3.0$, and the dotted line corresponds to $a = 10$.

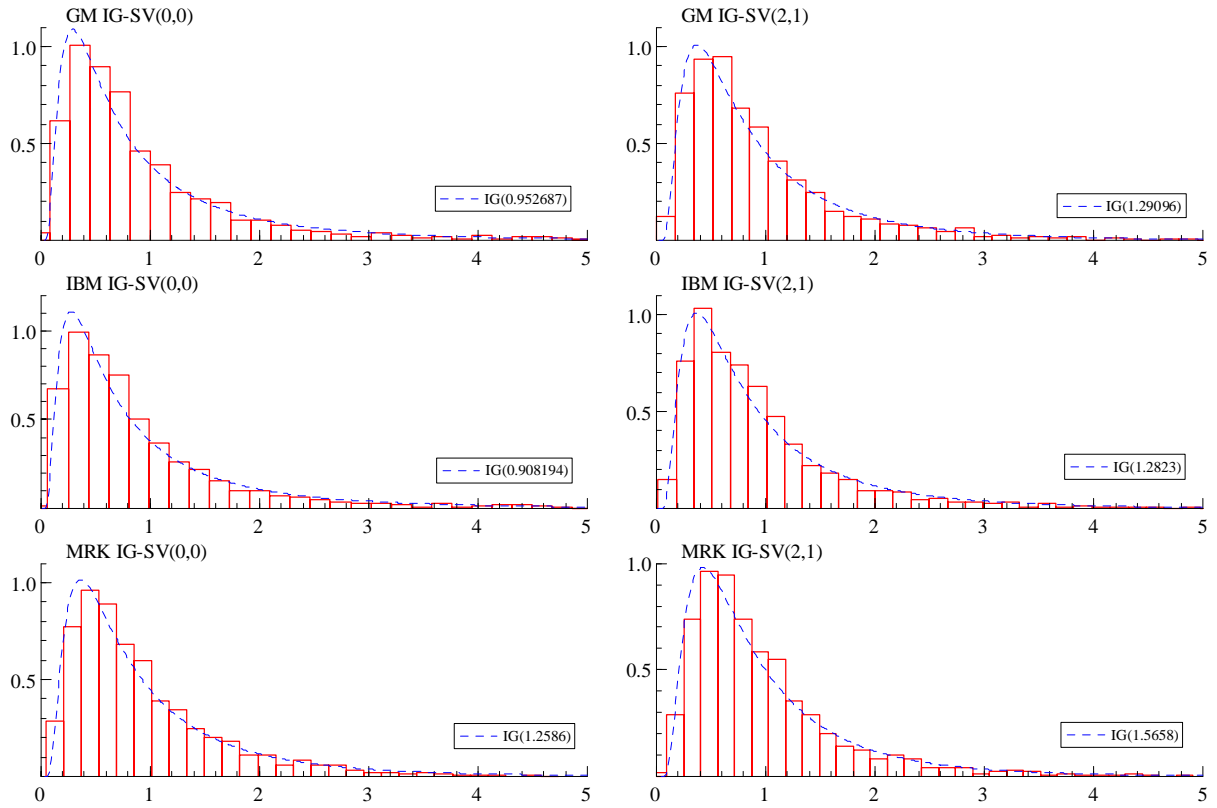


Figure 3: This figure shows the density of the RV_t with fitted IG density superimposed as a dotted line in the left hand plots. The right hand plots show the standardized RV_t , that is the estimated series, from the $IG-SV(2,1)$ specification again with the IG density as a dotted line. From top to bottom the plots are for GM, for IBM, and for MRK.

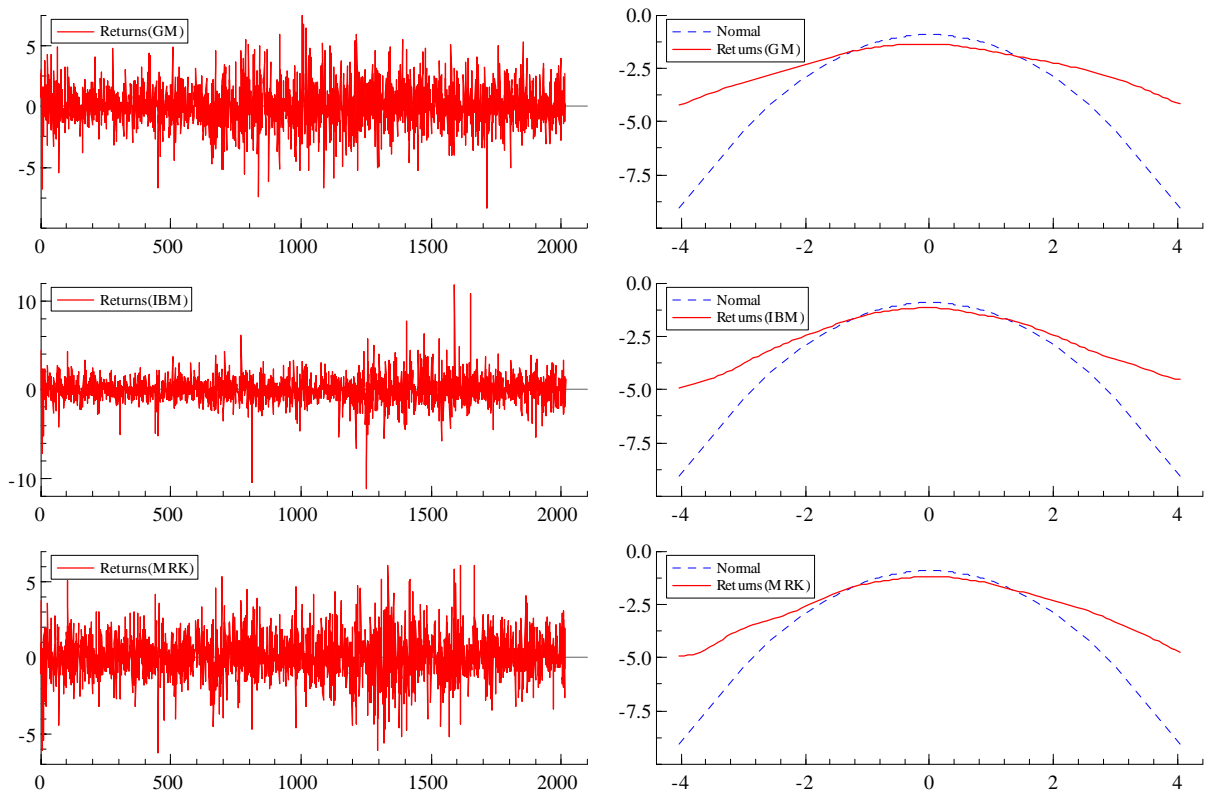


Figure 4: This figure shows time series plots of returns, R_t , in the left hand plots, and the log densities of the series in the right hand plot with the standardized Gaussian density superimposed as a dotted line. From top to bottom the plots are for GM, for IBM, and for MRK.




Figure 5: This plot shows the time series of standardized returns, $R_t/\sqrt{RV_t}$, in the left hand plots, and the log densities of the series in the right hand plot with the standardized Gaussian density superimposed as a dotted line. From top to bottom the plots are for GM, for IBM, and for MRK.

Table 1: Estimation results for IG-SV(p,q) models

Panel A: Estimation results for the IG-SV(0,0) model						
Underlying	GM	2014	IBM	2016	MRK	2013
<i>Likelihood</i>	-4114.70		-3503.96		-3338.91	
	Par.Estim.	Std.Error	Par.Estim.	Std.Error	Par.Estim.	Std.Error
a	0.9527	(0.0378)	0.9082	(0.0368)	1.2586	(0.0507)
	Statistic	P-value	Statistic	P-value	Statistic	P-value
$Q(20)$	891.185	[0.000]	774.213	[0.000]	432.160	[0.000]
$Q^2(20)$	147.807	[0.000]	68.683	[0.000]	29.297	[0.082]
$ARCH1 - 5$	13.110	[0.000]	7.118	[0.000]	0.099	[0.992]
<i>Schwarz</i>	4.086		3.476		3.317	
Panel B: Estimation results for the IG-SV(2,1) model						
Underlying	GM	2014	IBM	2016	MRK	2013
<i>Likelihood</i>	-3884.68		-3245.92		-3168.58	
	Par.Estim.	Std.Error	Par.Estim.	Std.Error	Par.Estim.	Std.Error
α_2	-0.0989	(0.0314)	-0.2317	(0.0359)	-0.1473	(0.0367)
α_1	0.1676	(0.0261)	0.2635	(0.0326)	0.2198	(0.0291)
β_1	0.9184	(0.0215)	0.9632	(0.0112)	0.9056	(0.0279)
a	1.2910	(0.0548)	1.2823	(0.0578)	1.5658	(0.0686)
	Statistic	P-value	Statistic	P-value	Statistic	P-value
$Q(20)$	16.593	[0.679]	14.475	[0.806]	15.107	[0.770]
$Q^2(20)$	7.157	[0.981]	8.536	[0.954]	8.684	[0.950]
$ARCH1 - 5$	0.448	[0.815]	0.096	[0.993]	0.168	[0.974]
<i>Schwarz</i>	3.859		3.222		3.149	

Notes: This table reports the estimation results for two models for RV. Panel A reports the results for the simple Constant Volatility case whereas Panel B reports the result of a more general IG-SV(2,1) specification for the dynamics. Thus, the latter model allows for time varying volatility. In parenthesis the standard errors are reported. $Q(20)$ is the Ljung-Box portmanteau test for up to 20'th order serial correlation in the standardized residuals, whereas $Q^2(20)$ is for up to 20'th order serial correlation in the squared standardized residuals. In square brakets next to all test statistics P-values are reported. The last row reports the Schwarz Information Criteria.

Table 2: American price estimates in a IG-SV option pricing model

Panel A: American put price estimates									
T	K	Gaussian-CV		IG-CV			IG-SV(2,1)		
		Price	Std.	Price	Std.	Rel Diff	Price	Std.	Rel Diff
7	85	0.000	(0.0001)	0.000	(0.0002)	-86.376%	0.001	(0.0003)	-94.259%
7	100	1.625	(0.0089)	1.603	(0.0119)	1.345%	1.591	(0.0119)	2.092%
7	115	15.000	(0.0000)	15.000	(0.0006)	-0.001%	15.001	(0.0021)	-0.004%
21	85	0.026	(0.0021)	0.029	(0.0022)	-10.224%	0.036	(0.0027)	-26.454%
21	100	2.760	(0.0147)	2.749	(0.0198)	0.395%	2.729	(0.0204)	1.112%
21	115	15.002	(0.0052)	15.003	(0.0062)	-0.005%	15.010	(0.0114)	-0.051%
63	85	0.463	(0.0119)	0.464	(0.0115)	-0.075%	0.474	(0.0122)	-2.279%
63	100	4.625	(0.0238)	4.616	(0.0309)	0.212%	4.584	(0.0319)	0.907%
63	115	15.391	(0.0349)	15.376	(0.0413)	0.097%	15.374	(0.0428)	0.110%
126	85	1.340	(0.0244)	1.337	(0.0219)	0.261%	1.339	(0.0232)	0.044%
126	100	6.324	(0.0268)	6.312	(0.0365)	0.188%	6.269	(0.0403)	0.882%
126	115	16.277	(0.0393)	16.273	(0.0514)	0.024%	16.242	(0.0515)	0.215%

Panel B: American call price estimates									
T	K	Gaussian-CV		IG-CV			IG-SV(2,1)		
		Price	Std.	Price	Std.	Rel Diff	Price	Std.	Rel Diff
7	115	0.001	(0.0003)	0.002	(0.0005)	-68.350%	0.003	(0.0007)	-81.918%
7	100	1.705	(0.0106)	1.684	(0.0121)	1.214%	1.672	(0.0122)	1.965%
7	85	15.066	(0.0068)	15.069	(0.0122)	-0.018%	15.078	(0.0125)	-0.082%
21	115	0.087	(0.0047)	0.092	(0.0054)	-6.024%	0.103	(0.0063)	-16.013%
21	100	2.994	(0.0219)	2.984	(0.0238)	0.326%	2.964	(0.0243)	1.000%
21	85	15.210	(0.0090)	15.212	(0.0186)	-0.014%	15.234	(0.0251)	-0.162%
63	115	0.993	(0.0229)	0.991	(0.0206)	0.137%	0.999	(0.0210)	-0.596%
63	100	5.314	(0.0398)	5.295	(0.0402)	0.352%	5.264	(0.0406)	0.946%
63	85	15.977	(0.0229)	15.972	(0.0394)	0.035%	15.980	(0.0425)	-0.016%
126	115	2.627	(0.0500)	2.623	(0.0435)	0.146%	2.618	(0.0460)	0.357%
126	100	7.656	(0.0657)	7.647	(0.0635)	0.110%	7.614	(0.0648)	0.550%
126	85	17.344	(0.0508)	17.339	(0.0596)	0.028%	17.337	(0.0624)	0.039%

Notes: This table shows American option prices for the IG-SV(p,q) option pricing models for a set of artificial options. The interest rate is fixed at 6% with a dividend yield of 3% both of which are annualized using 252 days a year. T denotes the time to maturity in days, K denotes the strike price, and for all the options the stock price is 100. The parameter values for the different models are the ones specified in the text. In the cross-sectional regressions powers of and cross-products between the stock level and the Realized Volatility level of total order less than or equal to three were used. Exercise is considered once every trading day. Prices reported are averages of 100 calculated prices using 20,000 paths and different seeds in the random number generator. In parentheses standard errors of these price estimates are reported. The columns headed Rel Diff reports the difference between the Gaussian-CV specification and corresponding IG-CV specification respectively the IG-SV(2,1) specification. Thus, they indicate the mispricing which would occur by the Gaussian-CV model if the true model is the corresponding IG model.

Table 3: Estimation results for NIG GARCH(p,q) models

Panel A: Estimation results for the NIG GARCH(0,0) model						
Underlying	GM	2014	IBM	2016	MRK	2013
<i>Likelihood</i>	-3977.19		-3585.11		-3609.26	
	Par.Estim.	Std.Error.	Par.Estim.	Std.Error.	Par.Estim.	Std.Error.
<i>a</i>	1.7330	(0.2994)	1.0109	(0.1348)	1.5795	(0.2446)
	Statistic	P-value	Statistic	P-value	Statistic	P-value
<i>Q</i> (20)	31.846	[0.045]	18.603	[0.548]	29.818	[0.073]
<i>Q</i> ² (20)	115.027	[0.000]	80.303	[0.000]	128.948	[0.000]
<i>ARCH</i> 1 – 5	7.167	[0.000]	10.973	[0.000]	6.117	[0.000]
<i>Schwarz</i>	3.950		3.557		3.586	
Panel B: Estimation results for the NIG GARCH(1,2) model						
Underlying	GM	2014	IBM	2016	MRK	2013
<i>Likelihood</i>	-3937.76		-3537.29		-3582.73	
	Par.Estim.	Std.Error.	Par.Estim.	Std.Error.	Par.Estim.	Std.Error.
β_1	0.9743	(0.0128)	0.9796	(0.0061)	0.9339	(0.0282)
α_2	-0.0234	(0.0231)	-0.1006	(0.0317)	-0.0424	(0.0316)
α_1	0.0453	(0.0207)	0.1168	(0.0312)	0.0798	(0.0292)
<i>a</i>	2.1980	(0.4266)	1.3224	(0.2051)	2.0509	(0.3748)
	Statistic	P-value	Statistic	P-value	Statistic	P-value
<i>Q</i> (20)	21.252	[0.382]	12.242	[0.907]	17.179	[0.641]
<i>Q</i> ² (20)	22.571	[0.164]	9.431	[0.926]	22.273	[0.174]
<i>ARCH</i> 1 – 5	0.616	[0.688]	0.861	[0.507]	0.412	[0.841]
<i>Schwarz</i>	3.912		3.511		3.561	

Notes: This table reports the estimation results for two GARCH models. Panel A reports the results for the simple Constant Volatility case whereas Panel B reports the result of a more general NIG GARCH(1,2) specification for the dynamics. Thus, the latter model allows for time varying volatility. In parenthesis the standard errors are reported. $Q(20)$ is the Ljung-Box portmanteau test for up to 20'th order serial correlation in the standardized residuals, whereas $Q^2(20)$ is for up to 20'th order serial correlation in the squared standardized residuals. In square brackets below all test statistics P-values are reported. The last row reports the Schwarz Information Criteria.

Table 4: American price estimates in a NIG-GARCH option pricing model

Panel A: American put price estimates									
T	K	Gaussian-CV		NIG-CV			NIG-GARCH(1,2)		
		Price	Std.	Price	Std.	Rel Diff	Price	Std.	Rel Diff
7	85	0.000	(0.0001)	0.000	(0.0001)	-81.615%	0.001	(0.0004)	-96.051%
7	100	1.625	(0.0089)	1.610	(0.0102)	0.935%	1.617	(0.0106)	0.467%
7	115	15.000	(0.0000)	15.000	(0.0003)	0.000%	15.000	(0.0014)	-0.002%
21	85	0.026	(0.0021)	0.028	(0.0022)	-6.822%	0.041	(0.0029)	-35.671%
21	100	2.760	(0.0147)	2.754	(0.0178)	0.188%	2.772	(0.0190)	-0.441%
21	115	15.002	(0.0052)	15.003	(0.0069)	-0.007%	15.007	(0.0101)	-0.030%
63	85	0.463	(0.0119)	0.463	(0.0126)	0.193%	0.496	(0.0137)	-6.663%
63	100	4.625	(0.0238)	4.618	(0.0273)	0.158%	4.632	(0.0280)	-0.146%
63	115	15.391	(0.0349)	15.380	(0.0445)	0.071%	15.404	(0.0440)	-0.082%
126	85	1.340	(0.0244)	1.337	(0.0198)	0.262%	1.366	(0.0204)	-1.896%
126	100	6.324	(0.0268)	6.319	(0.0314)	0.074%	6.322	(0.0339)	0.034%
126	115	16.277	(0.0393)	16.278	(0.0450)	-0.005%	16.286	(0.0480)	-0.055%

Panel B: American call price estimates									
T	K	Gaussian-CV		NIG-CV			NIG-GARCH(1,2)		
		Price	Std.	Price	Std.	Rel Diff	Price	Std.	Rel Diff
7	115	0.001	(0.0003)	0.001	(0.0004)	-63.623%	0.004	(0.0009)	-86.658%
7	100	1.705	(0.0106)	1.688	(0.0128)	0.968%	1.696	(0.0127)	0.525%
7	85	15.066	(0.0068)	15.067	(0.0105)	-0.007%	15.071	(0.0119)	-0.038%
21	115	0.087	(0.0047)	0.091	(0.0047)	-4.846%	0.115	(0.0056)	-24.252%
21	100	2.994	(0.0219)	2.990	(0.0257)	0.134%	3.008	(0.0261)	-0.490%
21	85	15.210	(0.0090)	15.208	(0.0172)	0.012%	15.228	(0.0221)	-0.123%
63	115	0.993	(0.0229)	0.992	(0.0210)	0.124%	1.037	(0.0230)	-4.285%
63	100	5.314	(0.0398)	5.302	(0.0419)	0.229%	5.318	(0.0439)	-0.083%
63	85	15.977	(0.0229)	15.972	(0.0334)	0.034%	16.004	(0.0356)	-0.165%
126	115	2.627	(0.0500)	2.621	(0.0418)	0.227%	2.653	(0.0435)	-0.976%
126	100	7.656	(0.0657)	7.644	(0.0581)	0.154%	7.650	(0.0590)	0.077%
126	85	17.344	(0.0508)	17.336	(0.0523)	0.044%	17.359	(0.0545)	-0.089%

Notes: This table shows American option prices for the NIG-GARCH option pricing models for a same set of artificial options as in Table 2. The parameter values for the different models are the ones specified in the text. The columns headed Rel Diff reports the difference between the Gaussian-CV specification and corresponding NIG-CV specification respectively the NIG-GARCH(1,2) specification. Thus, they indicate the mispricing which would occur by the Gaussian-CV model if the true model is the corresponding NIG model.

Table 5: Monte Carlo results on difference between using squared returns and Realized Volatility

Parameter	Initial Value	Squared Returns		Realized Volatility		Full Model	
		Par.Estim.	Std.Error.	Par.Estim.	Std.Error.	Par.Estim.	Std.Error.
β_1	0.93	0.9506	(0.0027)	0.9233	(0.0028)	0.9228	(0.0029)
$\alpha_{EiV,2}$	0.00	-0.0295	(0.0020)			-0.0010	(0.0017)
$\alpha_{EiV,1}$	0.00	0.0473	(0.0020)			-0.0001	(0.0017)
$\alpha_{RV,2}$	-0.16			-0.1546	(0.0037)	-0.1540	(0.0039)
$\alpha_{RV,1}$	0.21			0.2069	(0.0035)	0.2075	(0.0036)
a	1.40	1.3140	(0.0152)	1.4322	(0.0148)	1.4368	(0.0149)

Notes: This table reports results from the Monte Carlo study of the difference between using squared returns and Realized Volatility in estimating the variance equation. The tabel reports results from estimating NIG GARCH models on IG-SV data using the noisy daily squared returns, the actual RV series, and both as regressors.

Table 6: Estimation results for the NIG GARCH(1,2) model augmented with the observed RV series

	GM	2014	IBM	2016	MRK	2013
<i>Likelihood</i>		-3926.15		-3526.39		-3571.81
	Par.Estim.	Std.Error.	Par.Estim.	Std.Error.	Par.Estim.	Std.Error.
β	0.9446	(0.1452)	0.9738	(0.0137)	0.9477	(0.0336)
$\alpha_{EiV,2}$	0.0032	(0.0276)	-0.0605	(0.0320)	0.0013	(0.0298)
$\alpha_{EiV,1}$	-0.0017	(0.0251)	0.0672	(0.0320)	0.0207	(0.0299)
$\alpha_{RV,2}$	-0.0661	(0.1436)	-0.1840	(0.0630)	-0.2703	(0.0807)
$\alpha_{RV,1}$	0.1151	(0.0510)	0.2002	(0.0626)	0.2829	(0.0748)
a	2.3601	(0.4891)	1.3845	(0.2172)	2.2666	(0.4342)
	Statistic	P-value	Statistic	P-value	Statistic	P-value
$Q(20)$	21.527	[0.367]	12.389	[0.902]	16.252	[0.701]
$Q^2(20)$	18.456	[0.239]	12.766	[0.620]	24.383	[0.059]
$ARCH1 - 5$	0.206	[0.960]	1.222	[0.296]	0.546	[0.742]
<i>Schwarz</i>	3.901		3.500		3.551	

Notes: This table reports the estimation results for the NIG GARCH(1,2) model using the RV series in addition to the original specification denoted EiV. In parenthesis the standard errors are reported. $Q(20)$ is the Ljung-Box portmanteau test for up to 20'th order serial correlation in the standardized residuals, whereas $Q^2(20)$ is for up to 20'th order serial correlation in the squared standardized residuals. In square brakets below all test statistics P-values are reported. The last row reports the Schwarz Information Criteria.

Table 7: Overall performance for the IG-SV option pricing model

Panel A: GM						
Model	All	(1214)	Put	(350)	Call	(864)
	RBIAS	RSE	RBIAS	RSE	RBIAS	RSE
IG-SV(0,0)	14.58%	5.88%	15.34%	6.57%	14.27%	5.60%
IG-SV(1,0)	14.17%	5.60%	14.84%	6.00%	13.90%	5.43%
IG-SV(1,1)	5.35%	2.75%	5.22%	3.14%	5.40%	2.59%
IG-SV(2,1)	4.62%	2.60%	4.62%	3.01%	4.61%	2.43%
Best	4	4	4	4	4	4
Worst	1	1	1	1	1	1

Panel B: IBM						
Model	All	(4506)	Put	(1523)	Call	(2983)
	RBIAS	RSE	RBIAS	RSE	RBIAS	RSE
IG-SV(0,0)	-6.88%	8.69%	-10.11%	10.98%	-5.22%	7.52%
IG-SV(1,0)	-6.57%	8.13%	-9.84%	10.23%	-4.90%	7.06%
IG-SV(1,1)	-1.29%	5.45%	-3.78%	7.06%	-0.02%	4.62%
IG-SV(2,1)	0.76%	4.38%	-1.81%	5.53%	2.07%	3.79%
Best	4	4	4	4	3	4
Worst	1	1	1	1	1	1

Panel C: MRK						
Model	All	(1433)	Put	(246)	Call	(1187)
	RBIAS	RSE	RBIAS	RSE	RBIAS	RSE
IG-SV(0,0)	15.62%	9.04%	23.37%	16.45%	14.02%	7.50%
IG-SV(1,0)	15.21%	8.66%	22.70%	15.41%	13.66%	7.26%
IG-SV(1,1)	9.09%	5.91%	12.21%	9.02%	8.44%	5.26%
IG-SV(2,1)	6.60%	5.09%	8.49%	7.89%	6.21%	4.51%
Best	4	4	4	4	4	4
Worst	1	1	1	1	1	1

Notes: This table reports results on the overall pricing performance of the IG-SV option pricing model using the metrics from the text. In particular, letting P_k and \tilde{P}_k denote the k 'th observed price respectively the k 'th model price these are the relative mean bias, $RBIAS \equiv K^{-1} \sum_{k=1}^K \frac{(\tilde{P}_k - P_k)}{P_k}$, and the relative mean squared error, $RSE \equiv K^{-1} \sum_{k=1}^K \frac{(\tilde{P}_k - P_k)^2}{P_k^2}$.

Table 8: Overall performance for the IG-SV option pricing model in terms of moneyness

Panel A: GM						
Model	OTM	(547)	ATM	(318)	ITM	(349)
	RBIAS	RSE	RBIAS	RSE	RBIAS	RSE
IG-SV(0,0)	20.66%	10.40%	15.96%	4.10%	3.79%	0.41%
IG-SV(1,0)	20.31%	9.99%	15.22%	3.75%	3.61%	0.39%
IG-SV(1,1)	6.53%	4.88%	8.11%	1.84%	0.99%	0.24%
IG-SV(2,1)	5.46%	4.62%	7.37%	1.71%	0.79%	0.22%
Best	4	4	4	4	4	4
Worst	1	1	1	1	1	1

Panel B: IBM						
Model	OTM	(2212)	ATM	(939)	ITM	(1355)
	RBIAS	RSE	RBIAS	RSE	RBIAS	RSE
IG-SV(0,0)	-13.05%	15.79%	0.61%	3.64%	-1.99%	0.59%
IG-SV(1,0)	-12.17%	14.76%	-0.05%	3.43%	-1.96%	0.58%
IG-SV(1,1)	-3.91%	9.78%	3.98%	2.53%	-0.67%	0.40%
IG-SV(2,1)	-0.72%	7.80%	5.37%	2.23%	-0.02%	0.29%
Best	4	4	2	4	4	4
Worst	1	1	4	1	1	1

Panel C: MRK						
Model	OTM	(492)	ATM	(392)	ITM	(549)
	RBIAS	RSE	RBIAS	RSE	RBIAS	RSE
IG-SV(0,0)	27.25%	20.60%	16.96%	6.20%	4.25%	0.70%
IG-SV(1,0)	26.61%	19.72%	16.36%	5.94%	4.17%	0.69%
IG-SV(1,1)	14.84%	13.19%	10.82%	4.34%	2.70%	0.50%
IG-SV(2,1)	10.20%	11.38%	8.55%	3.72%	1.99%	0.43%
Best	4	4	4	4	4	4
Worst	1	1	1	1	1	1

Notes: This table reports results on the overall pricing performance in terms of moneyness of the IG-SV option pricing model using the metrics from the text. In particular, letting P_k and \tilde{P}_k denote the k 'th observed price respectively the k 'th model price these are the relative mean bias, $RBIAS \equiv K^{-1} \sum_{k=1}^K \frac{(\tilde{P}_k - P_k)}{P_k}$, and the relative mean squared error, $RSE \equiv K^{-1} \sum_{k=1}^K \frac{(\tilde{P}_k - P_k)^2}{P_k^2}$.

Table 9: Pricing performance for the IG-SV option pricing model relative to the NIG GARCH option pricing model

Panel A: GM						
Model	All	(1214)	Put	(350)	Call	(864)
	RBIAS	RSE	RBIAS	RSE	RBIAS	RSE
SV(0,0)/GARCH(0,0)	-0.020	-0.016	-0.022	-0.033	-0.020	-0.009
SV(1,0)/GARCH(0,1)	-0.031	-0.045	-0.032	-0.078	-0.030	-0.029
SV(1,1)/GARCH(1,1)	0.022	-0.076	-0.217	-0.194	0.162	-0.004
SV(2,1)/GARCH(1,2)	-0.153	-0.132	-0.328	-0.222	-0.053	-0.078
Max	0.022	-0.016	-0.022	-0.033	0.162	-0.004
Min	-0.153	-0.132	-0.328	-0.222	-0.053	-0.078

Panel B: IBM						
Model	All	(4506)	Put	(1523)	Call	(2983)
	RBIAS	RSE	RBIAS	RSE	RBIAS	RSE
SV(0,0)/GARCH(0,0)	-0.012	0.010	-0.008	0.011	-0.017	0.009
SV(1,0)/GARCH(0,1)	-0.103	-0.020	-0.092	-0.029	-0.114	-0.013
SV(1,1)/GARCH(1,1)	-0.897	-0.048	-0.782	-0.132	-0.998	0.030
SV(2,1)/GARCH(1,2)	-0.938	-0.274	-0.894	-0.349	-0.791	-0.207
Max	-0.012	0.010	-0.008	0.011	-0.017	0.030
Min	-0.938	-0.274	-0.894	-0.349	-0.998	-0.207

Panel C: MRK						
Model	All	(1433)	Put	(246)	Call	(1187)
	RBIAS	RSE	RBIAS	RSE	RBIAS	RSE
SV(0,0)/GARCH(0,0)	0.012	0.016	-0.004	0.022	0.018	0.013
SV(1,0)/GARCH(0,1)	0.003	0.000	-0.016	-0.008	0.010	0.003
SV(1,1)/GARCH(1,1)	-0.017	-0.013	-0.125	-0.020	0.021	-0.011
SV(2,1)/GARCH(1,2)	-0.252	-0.121	-0.361	-0.113	-0.214	-0.123
Max	0.012	0.016	-0.004	0.022	0.021	0.013
Min	-0.252	-0.121	-0.361	-0.113	-0.214	-0.123

Notes: This table reports results on the pricing performance of the IG-SV option pricing model relative to that of the NIG GARCH model. In the table a negative cell value indicates that the IG-SV specification has smaller pricing errors than a NIG GARCH model with equal number of parameters. A positive cell value indicates that the NIG GARCH model is the best performing model. The actual cell value indicates how much better or worse the IG-SV model performed relative to the NIG GARCH model.

Table 10: Pricing performance for the IG-SV option pricing model relative to the NIG GARCH option pricing model in terms of moneyness

Panel A: GM						
Model	OTM	(547)	ATM	(318)	ITM	(349)
	RBIAS	RSE	RBIAS	RSE	RBIAS	RSE
SV(0,0)/GARCH(0,0)	-0.015	-0.018	-0.022	-0.008	-0.057	-0.031
SV(1,0)/GARCH(0,1)	-0.019	-0.042	-0.045	-0.056	-0.079	-0.054
SV(1,1)/GARCH(1,1)	-0.042	-0.099	0.194	0.051	-0.255	-0.078
SV(2,1)/GARCH(1,2)	-0.238	-0.151	0.057	-0.029	-0.424	-0.116
Max	-0.015	-0.018	0.194	0.051	-0.057	-0.031
Min	-0.238	-0.151	-0.045	-0.056	-0.424	-0.116

Panel B: IBM						
Model	OTM	(2212)	ATM	(939)	ITM	(1355)
	RBIAS	RSE	RBIAS	RSE	RBIAS	RSE
SV(0,0)/GARCH(0,0)	-0.007	0.010	0.133	0.012	-0.040	0.012
SV(1,0)/GARCH(0,1)	-0.084	-0.022	-0.924	0.000	-0.130	-0.017
SV(1,1)/GARCH(1,1)	-0.815	-0.079	-0.256	0.370	-0.818	-0.050
SV(2,1)/GARCH(1,2)	-0.965	-0.299	0.025	0.051	-0.995	-0.356
Max	-0.007	0.010	0.133	0.370	-0.040	0.012
Min	-0.965	-0.299	-0.924	0.000	-0.995	-0.356

Panel C: MRK						
Model	OTM	(492)	ATM	(392)	ITM	(549)
	RBIAS	RSE	RBIAS	RSE	RBIAS	RSE
SV(0,0)/GARCH(0,0)	0.013	0.022	-0.001	-0.008	0.041	0.013
SV(1,0)/GARCH(0,1)	0.003	0.004	-0.008	-0.018	0.038	0.005
SV(1,1)/GARCH(1,1)	-0.028	-0.016	-0.021	-0.002	0.055	-0.012
SV(2,1)/GARCH(1,2)	-0.295	-0.120	-0.199	-0.121	-0.187	-0.134
Max	0.013	0.022	-0.001	-0.002	0.055	0.013
Min	-0.295	-0.120	-0.199	-0.121	-0.187	-0.134

Notes: This table reports results on the pricing performance of the IG-SV option pricing model relative to that of the NIG GARCH model in terms of moneyness. In the table a negative cell value indicates that the IG-SV specification has smaller pricing errors than a NIG GARCH model with equal number of parameters. A positive cell value indicates that the NIG GARCH model is the best performing model. The actual cell value indicates how much better or worse the IG-SV model performed relative to the NIG GARCH model.

Research Papers 2008



- 2008-01 John A. Carlson, Christian M. Dahl, and Carol L. Osler: Short-run Exchange-Rate Dynamics: Theory and Evidence
- 2008-02 Peter Reinhard Hansen: Reduced-Rank Regression: A Useful Determinant Identity
- 2008-03 Søren Johansen, Katarina Juselius, Roman Frydman and Michael Goldberg: Testing hypotheses in an I(2) model with applications to the persistent long swings in the Dmk/\$ rate
- 2008-04 Olaf Posch: Explaining output volatility: The case of taxation
- 2008-05 Annastiina Silvennoinen and Timo Teräsvirta: Modelling Multivariate Autoregressive Conditional Heteroskedasticity with the Double Smooth Transition Conditional Correlation GARCH Model
- 2008-06 Annastiina Silvennoinen and Timo Teräsvirta: Multivariate GARCH models. To appear in T. G. Andersen, R. A. Davis, J.-P. Kreiss and T. Mikosch, eds. Handbook of Financial Time Series. New York: Springer.
- 2008-07 Changli He, Annastiina Silvennoinen and Timo Teräsvirta: Parameterizing unconditional skewness in models for financial time series
- 2008-08 Cristina Amado and Timo Teräsvirta: Modelling Conditional and Unconditional Heteroskedasticity with Smoothly Time-Varying Structure
- 2008-09 Søren Johansen and Bent Nielsen: An analysis of the indicator saturation estimator as a robust regression estimator
- 2008-10 Peter Christoffersen, Kris Jacobs, Christian Dorion and Yintian Wang: Volatility Components, Affine Restrictions and Non-Normal Innovations
- 2008-11 Peter Christoffersen, Kris Jacobs, Chayawat Ornthanalai and Yintian Wang: Option Valuation with Long-run and Short-run Volatility Components
- 2008-12 Tom Engsted and Stig V. Møller: An iterated GMM procedure for estimating the Campbell-Cochrane habit formation model, with an application to Danish stock and bond returns
- 2008-13 Lars Stentoft: Option Pricing using Realized Volatility