

This is an Open Access document downloaded from ORCA, Cardiff University's institutional repository: <https://orca.cardiff.ac.uk/id/eprint/127453/>

This is the author's version of a work that was submitted to / accepted for publication.

Citation for final published version:

Calice, Giovanni, Chen, Jing and Williams, Julien 2020. Forecasting options prices using discrete time volatility models estimated at mixed timescales. *Journal of Derivatives* 27 (3) , pp. 45-74. 10.3905/jod.2019.1.094

Publishers page: <https://doi.org/10.3905/jod.2019.1.094>

Please note:

Changes made as a result of publishing processes such as copy-editing, formatting and page numbers may not be reflected in this version. For the definitive version of this publication, please refer to the published source. You are advised to consult the publisher's version if you wish to cite this paper.

This version is being made available in accordance with publisher policies. See <http://orca.cf.ac.uk/policies.html> for usage policies. Copyright and moral rights for publications made available in ORCA are retained by the copyright holders.



Forecasting Options Prices Using Discrete Time Volatility Models Estimated at Mixed Timescales

GIOVANNI CALICE, JING CHEN, AND JULIAN WILLIAMS

GIOVANNI CALICE

is a senior lecturer in Finance at the School of Business and Economics at Loughborough University, UK.
g.calice@lboro.ac.uk

JING CHEN

is a reader in financial mathematics at the School of Mathematics at Cardiff University, UK.
chenj60@cardiff.ac.uk

JULIAN WILLIAMS

is a professor of finance at Durham University Business School and Executive Director of the Institute of Hazard Risk and Resilience at Durham University, UK.
julian.williams@durham.ac.uk

*All articles are now categorized by topics and subtopics. [View at PM-Research.com.](https://www.pm-research.com)

KEY FINDINGS

- In this paper, we present a new method for estimating the parameters for a jump GARCH model. We provide a series of empirical tests of the efficacy of the GARCH type option models. We analyse the S&P 500 index and for a sample of 20 individual equities sampled from the Dow Jones 30. Our out-of-sample test covers over a third of a million individually equity traded prices.
- We find three primary empirical results. First, pre-filtering for jumps improves the accuracy of options models based on GARCH processes. Second, for certain stocks, models that explicitly incorporate jumps substantially outperform all other models. Third, for the S&P500, the GARCH model estimated on jump-filtered returns appears to dominate.

ABSTRACT: *Option pricing models traditionally have utilized continuous-time frameworks to derive solutions or Monte Carlo schemes to price the contingent claim. Typically these models were calibrated to discrete-time data using a variety of approaches. Recent work on GARCH based option pricing models have introduced a set of models that easily can be estimated via MLE or GMM directly from discrete time spot data. This paper provides a series of extensions to the standard discrete-time options pricing setup and then implements a set of various pricing approaches for a very large cross-section of equity and index options against the forward-looking traded market price of these options, out-of-sample. Our analysis provides two significant findings. First, we provide clear evidence that including autoregressive jumps in the options model is critical in determining the correct price of heavily out-of-the money and in-the-money options relatively close to maturity. Second, for longer maturity options, we show that the anticipated performance of the popular component*

GARCH models, which exhibit long persistence in volatility, does not materialize. We ascribe this result in part to the inherent instability of the numerical solution to the option price in the presence of component volatility. Taken together, our results suggest that when pricing options, the first best approach is to include jumps directly in the model, preferably using jumps calibrated from intraday data.

TOPICS: *Options, volatility measures**

Option pricing on equities and equity indexes has been the subject of extended discussion since Black, Scholes and Merton provided their foundational insights. The development of this research area has focused on continuous-time models and tractable solutions that are either set around the log-normal pricing paradigm or use stochastic volatility potentially in the presence of jumps.

In discussing the 1987 crash, Bates (1996) noted that in the analysis of options without the formal empirical and quantitative analysis of both the historical data of the underlying market and forward-looking calibration from the options market, complete insight into either market would prove elusive. Part of our contribution is to add another dimension to this context. Asset pricing in both markets is also dependent on the analysis of multiple timescales. We demonstrate that when including discontinuous jumps and the use of at least two timescales, the extra information provided is more important to forecasting the value of options than to improving the maximum likelihood estimation of the underlying data-generating process.

In a separate tract of finance research, in discrete time-series econometrics, the generalized autoregressive conditionally heteroskedastic (GARCH) model has proved extremely popular for a variety of applications related to forecasting volatility in discrete time. However, the consistent composition of the two models in an analytic framework has been a more recent phenomenon. In an important series of contributions, Bates (1996), Heston and Nandi (2000), and Christoffersen et al. (2008), among others, have provided a significant step forward in tractable options pricing not seen since the derivation of the original Black and Scholes models and their subsequent extensions in the early and mid-1970s. The key insight of these seminal papers is to combine discrete-time models of volatility and jumps with risk-neutral dynamics.

Most strikingly, the results of Christoffersen et al. (2008) suggested that for S&P 500 index options, a highly persistent form of GARCH model, the CGARCH, provides the best out-of-sample performance when pricing contracts against those reported in the market, for all maturities and money-ness. This was especially important as the comparison models included an autoregressive Poisson point process driving the jump diffusion. This result goes against much of the preceding literature and industry practice for which Poisson type jump processes are considered essential, in particular for the pricing of near-maturity options, which are not at- or near-the-money.

This paper makes three main contributions to the literature. First, we provide a series of empirical tests of the efficacy of the GARCH type option models suggested in Christoffersen et al. (2008) out-of-sample for both the S&P 500 index and for a sample of 20 individual equities sampled from the DOW Jones 30. Our major

contribution is to propose alternative estimation procedures over the domain of models, and provide evidence from a far larger pool of equities and equity indexes for this analysis. Our primary argument is that jumps are an intraday feature and that maximum likelihood evidence at a daily frequency, commonly used in option pricing, inefficiently detects the intraday jump phenomena. Using a heterogeneous timescale approach should yield a better fit when comparing estimated versus market prices of options. A key benefit of this strategy is that it avoids the problems identified in Durham et al. (2015), who note that the standard maximum likelihood estimator of the unobserved volatility and jump processes does not meet the normal requirements of a standard filter for an AR-Jump model. In this case the jumps are partially observed, such that the existence of $N_t > 0$ is imputed from high frequency data at an a-priori stage.

Second, our out-of-sample test covers over a third of a million individually traded prices sampled from the Options Price Reporting Authority (OPRA) data feed for 2012 and 2013, and this is, to our knowledge, the largest data set ever constructed. Third, we implement an original estimation method for the jump intensity process that uses data at both the daily and intraday frequencies. In summary, we believe that this is the most comprehensive study of its type ever undertaken, and that it should provide a reasonably definitive conclusion to the debate on the performance of this model.

Our results indicate two important facets. First, if the option pricing model does not, or cannot, explicitly incorporate jumps in its estimation stage, then pre-filtering for jumps improves the accuracy of options models based around standard GARCH processes. Second, for certain stocks, such as IBM, models that explicitly incorporate jumps substantially outperform all other models. However, for the S&P 500 index, the results are less conclusive. Here, the standard GARCH model of Heston and Nandi (2000) outperforms all others out-of-sample for both root mean square error (RMSE) and median absolute error (MAE). This level of accuracy is in line with that found for similar modeling fitting against actual options data, such as that found in Ballotta et al. (2017).

The rest of this paper is organized as follows: The next section briefly reviews related studies and highlights our contribution to the existing literature. Then we describe our methodology and the rationale regarding the key assumptions in our forecasting and estimation procedures. Then we describe our rich data set, and

then present and discuss the empirical results. The final section concludes with some suggestions for extended work in this area.

MOTIVATION AND RELATED LITERATURE

Jumps are often considered primarily to be a short term phenomenon, and prior empirical studies focused on near maturity, near the money options. Our dataset encompasses a broad range of maturities and we show that the relative effectiveness of different option pricing tools is significantly affected by moneyness and maturity.

Indeed, several studies have argued that jumps either do not exist or are being spuriously detected; see for instance Bajgrowicz and Scaillet (2009) who utilize a short sample of data for the Dow Jones constituents and argue that sudden price movements are more likely the result of a “sudden burst of volatility.” How sudden a “burst” is required to shift the likelihood of a 3% jump from effectively zero to 0.02% in a day? It needs to be a burst that amounts to reducing the number of standard deviations between the current price and the out-of-the-money threshold from 364 to 2.87. The answer is just a little above 1,000% annualized; when a log-normal distribution approaches this level of variance the mode of the distribution is indistinguishable from zero.

At a lower-frequency modeling, clustering in volatility and sudden discontinuous jumps in returns has been a topic of substantial interest since the inception of the financial economics literature. Motivated by the need for accurate estimation of volatility to price derivatives contracts, Roberts (1959) and Merton (1976) were among the first to discuss the non-Gaussian properties of returns in relation to derivatives contracts.

The ARCH/GARCH models of Engle (1982) and Bollerslev (1986) were the first to explicitly capture autoregressive clustering of volatility in a time series framework. The standard form of the GARCH models and most of the endless contemporary variations do not provide simple closed-form solutions to the price of a contingent claim on an asset exhibiting return behavior with these effects. However, as previously noted, several recent contributions have derived closed-form solutions; two of the most prominent are those of Heston and Nandi (2000) and Christoffersen et al. (2008).

The general consensus in the mainstream option pricing literature is to assume that the GARCH volatility component is a continuous diffusion component

of returns and all other càdlàg components are referred to as discontinuous jumps. A useful summary of the current state of art on discretely sampling continuous diffusions can be found in Mykland and Zhang (2009). An important feature of the literature on modeling volatility is that while unbiased estimates of volatility are possible for most types of continuous diffusions with fixed time horizons T , given increasing sampling frequency, this is not true for the drift.

A recent strand of the literature has focused on extracting the discontinuous jump component from the continuous diffusion. This is motivated in part by versions of the example we described above. When options are traded near maturity, any price that deviates substantively from $|S_t - K|$ must be driven by some very short term time component. Over longer maturities this is less important. However, for continuous trading, small differences in pricing of the contingent claim can be extremely critical. It should also be noted that, even ex-post maturity, there is no definitive “correct” price for a contingent claim of this type. Therefore, the purpose of an option pricing model is to correctly predict the market price of an option and not $|S_T - K|$.

On this note, Andersen et al. (2002) have documented empirically that the inclusion of jumps reduces systematic bias in continuous-time stochastic volatility models based around the specification first presented in Heston (1993). The magnitude of the results in Andersen et al. (2002) are broadly representative of the improvement in fit when including the jump process in the underlying diffusion. A non-exhaustive list of examples of similar results, for both discrete- and continuous-time models, can be found in Chan and Maheu (2002); Jorion (1998); Maheu and McCurdy (2004), and Palm and Vlaar (1993).

More specific forms of jumps, for instance Poisson-normal jumps, have been found to improve the fit from the historical evolution of the underlying asset process to that inferred by the options market. See for instance Bakshi et al. (1997), Bates (1996, 2006), Chernov et al. (2003), Eraker et al. (2003), Eraker (2004), and Pan (2002), among others. In fact there is no empirical evidence to indicate that the exclusion of a jump process from the underlying can actually improve the fit of the predicted option price to the prices observed in the market.

Pricing options tends to be primarily from low frequency (circa daily). However, jumps generally

are thought to be a “high-frequency” phenomenon. Aït-Sahalia (2004) offers a comprehensive summary on jump diffusion models and their applications to asset pricing in addition to a substantial discussion on how to filter the jump component from the continuous diffusion from high-frequency data. The motivation for extracting the jump component primarily lies in the impact of jumps on portfolio risk management and derivative pricing. The presence of jumps requires a substantial adjustment to derivative pricing models. In particular, the dynamics of the intensities (a latent variable) of a jump component present a variety of challenges. Aït-Sahalia (2004) develops a method for extracting the quadratic variation of the continuous diffusion for the two common forms of jumps, infinitely active Lévy and finitely active Poisson.

Other empirical work in this area has focused on extracting finitely active jumps—see, for example, Andersen et al. (2003), Barndorff-Nielson and Shephard (2004), Barndorff-Nielson and Shephard (2006), Tauchen and Zhou (2011), and more recently Fuertes and Olmo (2012). Visual inspection of the return series suggests that during the sample period some jumps have occurred. However, these discontinuities are mostly visible at the higher sampling frequencies. Option pricing models that contain persistent volatility or jump effects, such as that demonstrated in Christoffersen et al. (2008), are usually calibrated from daily data. However, a growing body of the econometric literature (see for instance Tauchen and Zhou, 2011) suggest that filtering for jumps at relatively low frequencies (such as daily or weekly data) is problematic, if not unfeasible. Therefore a model estimated only on daily returns might well be inconsistently and/or inaccurately calibrated if only fitted to low frequency data.

Christoffersen and Jacobs (2004) find that for option valuation purposes the best fitting model is a parsimonious GARCH model that allows only for volatility clustering and standard leverage effects.

Corsi et al. (2013) develop a reduced form discrete-time stochastic volatility option pricing model that exploits the information of high-frequency data, which is the measure of realized volatility. By modeling the conditional mean of the volatility process using the Heterogeneous Autoregressive (HAR) multi-components model, they show, through an application to the SP500 index options, that this model outperforms competing GARCH-type and other stochastic volatility option pricing models.

Babaoglu et al. (2017) introduce a new class of models that incorporate three features: multiple volatility components, fat-tailed return innovations, and a variance-dependent price kernel. They apply these models on the SP500 option prices and find economically and statistically significant improvements over the benchmark SV option pricing model.

Overall, in the literature there is not a clear-cut consensus on the importance of jumps in options pricing. For example, Bakshi et al. (1997) demonstrate substantial benefits from including jumps in prices, whereas Bates (2000) find that such benefits are economically small, if not negligible. Furthermore, while studies using the time series of returns unanimously support jumps in prices, they disagree with respect to the importance of jumps in volatility. One plausible explanation for the above disparities is that most papers use data covering only short time periods. Since jumps are rare, short samples are likely to either over- or under-represent jumps and/or periods of high or low volatility, and thus could generate disparate results.

Broadie et al. (2007) use an extensive data set of S&P 500 futures options from January 1987 to March 2003 and find that adding price jumps to a square-root stochastic volatility (SV) model improves the cross-sectional fit by almost 50%. This is consistent with the large impact reported in Bakshi et al. (1997), but contrasts with the negligible gains documented in Bates (2000), Pan (2002), and Eraker (2004). Without any risk premium constraints, the SVJ and SVCJ models perform similarly in and out of sample. This is not surprising, as price jumps, which generate significant amounts of skewness and kurtosis, and stochastic volatility are clearly the two most important components for describing the time series of returns or for pricing options. Aït-Sahalia (2004), Carr and Wu (2003), and Huang and Tauchen (2005) also find evidence for jumps in prices.

Tauchen and Todorov (2010) have proposed a new continuous process to describe activity levels, which essentially extends the Blumenthal-Gettoor index to Ito semi-martingales. This study focuses on the activity signature function and suggests a more effective way of making non-parametric inference for generalized activities with finite sample but removing bias documented in the literature (see Aït-Sahalia and Jacod 2009a and (2009b, and Mykland et al. 2005). The model setups compare a pure-jump model to a pure-jump model with continuous component, and shows that the latter appears

to be the better model, for example, for examining jump activities in financial series.

Jacod and Todorov (2009) propose a bivariate discrete model to detect the occurrence of common jumps (at least one jump found simultaneously at the same time) or disjoint jumps. This study contributes to the literature by addressing the issue of existence of systematic jumps. Such models are not only theoretically important because of their simple and novel structure, but also practically useful. For example, at the aggregated market level, such a model can be used to explain whether multiple asset prices jump simultaneously, which forms a more systematic pattern, or vice versa, remain primarily an idiosyncratic phenomenon. The simulation methods prove that the process works for both infinite and some finite samples. In conclusion, these papers agree that diffusive stochastic volatility and jumps in prices are important, but they disagree over the importance of jumps in volatility.

Part of our approach is, therefore, to extract the jump component from the high-frequency data using the method of Tauchen and Zhou (2011) and then separately to estimate the daily volatility using the canonical GARCH type frameworks, namely the Component GARCH of Engle and Lee (1999) and the GARCH model with jumps, as proposed by Christoffersen et al. (2006) and Christoffersen et al. (2008), derived from the setting introduced by Heston and Nandi (2000).

When comparing option prices predicted from an underlying time series to market prices of options, the choice of the loss function is relatively critical; this is discussed at length in Christoffersen and Jacobs (2004). However, the major choice of the loss function advocated in Christoffersen and Jacobs (2004), the RMSE, contrasts with much of the contemporary literature on out-of-sample loss forecasting; see for instance Giacomini and Rossi (2010). While the RMSE applied directly to the actual and forecasted prices provides a useful “dollar” amount for the average pricing error, the option has a lower bound at zero. Hence, overpricing the option can be more significantly penalized than underpricing. This is particularly acute for heavily out of the money options, with prices intrinsically near zero. We therefore compute the RMSE on the log price of the predicted and realized option. We also provide RMSE error tables in the Online Supplement.

METHODOLOGY

We use intraday data to detect and extract jumps from returns. We follow a three-step procedure. We start by fitting an AR process to the arrival intensities of the AR-Jump process. In a second step, we fit a GARCH model to the jump-filtered returns. Finally, we impose the jump-filtered GARCH parameters on an AR-GARCH process and impute the AR process from the filtered returns. Against this setup, we develop a series of base cases. First, we estimate the AR-Jump-GARCH model fitted via maximum likelihood estimation at the daily timescale. Second, we estimate the AR-Jump-GARCH model with GARCH parameters imposed *after* we have extracted the jumps and the AR-jump model then fitted using maximum likelihood estimation. Finally, we compare these fits to a Component GARCH (CGARCH) model (the best-performing model in Christoffersen et al. 2008) and a basic GARCH specification. In each instance we compare in-sample fits via the fitted log-likelihoods, and most importantly we document how each model performs out-of-sample in the forward pricing of options.

The Price Process Assumptions

The basic price process of interest is the natural logarithm of the quoted price denoted $s(\tau)$, where $t_0 < \tau < T$, which is assumed to be a diffusion that exhibits varying volatility at heterogeneous timescales and potentially exhibiting discontinuous jumps. The objective of this paper is to infer the optimal approach to fitting the process to data and evaluate its performance in pricing stock options. The basic process of interest is as follows:

$$ds(\tau) = \mu_t d\tau + \sigma_t dW(\tau) + J_t dN(\tau), \quad (1)$$

where t is a daily index, μ_t is a daily level of drift, σ_t is a daily level of volatility, assumed to be of the form $\sigma_t = \sqrt{h_t}$ where h_t is a daily conditional variance and $J_t \sim \mathcal{N}(\mu_j, \sigma_j^2)$ is a random jump size. The random processes $W(\tau)$ and $N(\tau)$ are respectively a Wiener process and a Poisson point process such that $W(\tau + \Delta) - W(\tau) \sim \mathcal{N}(0, \Delta)$ and $\mathbb{P}(N(\tau + \Delta) - N(\tau) = n) = e^{-\tilde{\lambda}_t \Delta} (\tilde{\lambda}_t \Delta)^n / n!$, where $\tilde{\lambda}_t$ is the arrival intensity of the Poisson process and n is an integer.

For option pricing purposes, we assume that μ_t is determined by the model characteristics. Therefore, the key objects of interest are σ_t and $\tilde{\chi}_t$. We assume a variety of models for σ_t and $\tilde{\chi}_t$. Our basic specification is that σ_t is a GARCH model and $\tilde{\chi}_t$ is a stationary autoregressive process (AR-J). For our basic comparison, we follow Christoffersen et al. (2008) and compare a persistent, component GARCH model (denoted GARCH) and an AR-J model with a volatility mean-reverting GARCH model. Our approach is to fit the models in two stages. First, we extract the jumps by analyzing intra-daily data using a standard jump detection technique. We then fit the GARCH model to the de-jumped data at the daily frequency. For the final part of our analysis, we utilize the fitted models to price options and compare the goodness of fit to real options data via a proportional RMSE loss function.

Jump Detection and Extraction

Let us consider a high-frequency uniformly sampled grid indexed by i . The price diffusion is assumed to be sampled at $\tau_i \in \{\tau_0, \dots, \tau_n\}$, where $\tau_{i+\Delta} - \tau_i, \forall i \in \{0, \dots, n-1\}$ is constant over a day again indexed by t . Therefore, the time interval τ_0 to $\tau_n - 1$ is a single day of trading for the asset with log prices discretely recorded by $S_{t,i}$. We assume that $t \in \{1, \dots, T\}$ is a single year of data, usually ≈ 252 trading days.

Our objective is to separate $\int_{\tau_0}^{\tau_n} \sigma_s^2 ds$ from $\int_{\tau_0}^{\tau_n} (\sigma_s^2 + \lambda_s \sigma_j^2) ds$ and hence identify the jump and continuous diffusion components of the daily return $s(\tau_n) - s(\tau_0)$. Let $s_{t,i}$ be the log price for a regularly spaced grid at time index τ_i . If the process describing $s(\tau)$ is as that given in (1) then $RV_t = \sum_{i=1}^n r_{t,i}^2 \rightarrow \int_{\tau_0}^{\tau_n} (\sigma_s^2 + \lambda_s \sigma_j^2) ds$, where $r_{t,i} = S_{t,i} - S_{t,i-1}$. We assume that the contribution of the drift term μ_t over the course of a day is negligible and hence tends to zero. Following Barndorff-Nielson and Shephard (2004) and Tauchen and Zhou (2011), we define an alternative estimator, the bipower variation, that converges to $\int_{\tau_0}^{\tau_n} \sigma_s^2 ds$.

The bipower variation, BV_t is a member of a family of jump robust power estimators and has proven popular for jump detection approaches. This is defined as $BV_t = \frac{\pi}{2} \frac{n}{n-1} \sum_{i=2}^n |r_{t,i-1}| |r_{t,i}|$. In the absence of Poisson jumps $BV_t \rightarrow \tilde{\sigma}_t^2$. However, in the presence of jumps $\tilde{\sigma}_t^2 - BV_t \rightarrow \int_{\tau_0}^{\tau_n} \lambda_s \sigma_j^2 ds$ and the variation of

$\tilde{Z}_t = \tilde{\sigma}_t^2 - BV_t$ is asymptotically normal. As such, we can construct a standardized test such that

$$Z_t = \tilde{Z}_t \tilde{\sigma}_t^{-2} ((\pi^2/4 + \pi - 5) \frac{1}{n} \max[1, TQ_t/BV_t])^{-1/2}, \quad Z_t \sim \mathcal{N}(0,1)$$

with estimated jump size $\hat{J}_t = \text{sgn}(\tilde{R}_t)(RV_t - BV_t)^{1/2} \mathbf{1}_{\mathcal{N}(Z_t) \geq \mathcal{N}(Z^*)}$, where sgn is the signum function, $\mathbf{1}$ is the indicator function, Z^* is a standardized critical z-score, and $\mathcal{N}(Z)$ is the cumulative distribution function for a normal distribution with lower integrator z and $\tilde{R}_t = \sum_{i=1}^n r_{t,i}$. The term TQ_t is the without-jumps estimate of the quarticity of the continuous diffusion and is estimated using the tripower quarticity, suggested in Barndorff-Nielson and Shephard (2004) and Tauchen and Zhou (2011), among others. In this case, $TQ_t \rightarrow \int_{\tau_0}^{\tau_n} \sigma_s^4 ds \approx n \mu_{k=4/3} \frac{n}{n-2} \sum_{i=3}^n |r_{t,i-2}|^{4/3} |r_{t,i-1}|^{4/3} |r_{t,i}|^{4/3}$ where $\mu_k = 2^{k/2} \Gamma((k+1)/2) / \Gamma(k/2)$ and $\Gamma(k)$ denotes the Euler's Gamma function. Let $\tilde{N}_t = \mathbf{1}_{\mathcal{N}(Z_t) \geq \mathcal{N}(Z^*)}$, be the $\{0, 1\}$ count of at least one jump on day t . An issue that arises is the problem of $N_t > 1$ jumps. Note that the implicit assumption in Barndorff-Nielson and Shephard (2004) and the explicit assumption in Tauchen and Zhou (2011) is that $N_t = \{0, 1\}$. If we assume explicitly that the count process is Poisson then $\mathbb{P}(N_t > 1) > 0$, by construction. A simple approach is to assume that on a daily basis $\tilde{\chi}_t$ will be sufficiently small that the at-most one jump restriction does not cause substantial issues to the identification of $\tilde{\chi}_t$ from the time series of N_t . However, two alternative approaches exist. A first method is to sub-divide a day into blocks and sequentially test for jumps within those blocks. However, as the day is chopped into ever smaller blocks of time, the sample gets shorter for each block. A possible solution would be to shift to a higher sampling frequency, but this gives rise to potentially misleading issues such as stale prices. A second procedure is to compute J_t on the basis of the jump existence and then use a bootstrap resampling with replacement and robustly estimate μ_t and σ_j . For our pricing purposes, we will utilize an iterative approach based on maximum likelihood scores for our assumed jump process with GARCH.

We define a semi-parametric estimator for the count AR process in the spirit of Harvey and Fernandes (1989) and Heinen (2003). This is effectively an autoregressive conditional Poisson model (ACP) of the form

$$\tilde{\chi}_t = \tilde{\chi}_0 + \tilde{\rho} \tilde{\chi}_{t-1} + \tilde{\xi} \tilde{\eta}_{t-1}, \quad (2)$$

where the $\tilde{\theta}$ symbol for the parameter θ is used to distinguish the parameters and fitted intensities $\tilde{\chi}_t$ of the AR process from the nonparametric intraday jump detection analysis. Recall that for a Poisson process $\eta_t = N_t - \tilde{\chi}_t$ the unconditional mean, variance, and autocorrelation of this process are $\mathbb{E}(\tilde{N}_t) = \tilde{\chi} = \tilde{\chi}_0 / (1 - \tilde{\rho})$, $\text{var}(\tilde{N}_t) = \mathbb{E}(\tilde{N}_t - \tilde{\chi}) = \tilde{\chi}$ and $\text{corr}(\tilde{N}_t, \tilde{N}_{t-s}) = \tilde{\rho}^{s-1} \tilde{\rho} / (1 - \tilde{\rho}^2)$ respectively. This leaves us with two parameters, $\tilde{\rho}$ and $\tilde{\chi}_0$, that we can directly observe from the theoretical moments of the process and $\tilde{\xi}$ which we can solve from the filtration where $\mathbb{E}(\tilde{N}_t) = \tilde{\chi}_t$ and $\tilde{\xi}^2 \rightarrow T^{-1}(\tilde{N}_t - \tilde{\chi}_t)^2$ given the variance constraint in a second step. We can over-identify the parameters by computing $s > 1, \forall s \in \mathbb{N}_+$ autocorrelations. We then have two choices. We can embed the jumps in mean only or alternatively we can embed the jumps in both mean and variance as in Christoffersen et al. (2008). For our main specification we will embed the jumps in both mean and variance.

Using Mixed Frequency Data

Multiscale volatility estimation has received considerable attention in the literature recently. However, this has mostly focused on measuring ex-post quadratic variation. In order to price options under specific modeling assumptions, we need to utilize the multi-timescale data in a slightly different manner.

A simple iterative procedure for updating the counts and hence extracting the jump intensities is as follows:

Step 0: For each observation, compute the “jump free” return $R_t = \tilde{R}_t - \hat{J}_t$ and fit the GARCH model of your choice. Using the model

$$R_{t+1} = r + \lambda_z^{(J)} \hat{h}_t^{(J)}, \quad \hat{h}_t^{(J)} \sim \text{GARCH},$$

the coefficients from this pre-filtered GARCH in mean model also can be used as a benchmark against the unfiltered returns.

Step 1: Compute the initial unconditional mean and standard deviation from the detected jumps:

$$\begin{aligned} \tilde{\mu}_j^{(1)} &= \left(\sum_{t=1}^T \tilde{N}_t \right)^{-1} \sum_{t=1}^T \hat{J}_t, \\ \tilde{\sigma}_j^{(1)} &= \sqrt{\left(\sum_{t=1}^T \tilde{N}_t \right)^{-1} \sum_{t=1}^T \hat{J}_t - \mu_j^{(1)}}. \end{aligned}$$

This follows the Tauchen and Zhou (2011) assumption that only one jump has occurred within a day. It is possible to simply threshold these jumps and ascribe extra jumps to outliers. However, we can improve the estimation accuracy of detecting multiple jumps by computing the conditional jump intensity and commuting the overall likelihood of the process for one potentially additional jump.

Step 2: Set the initial count to unity, $\tilde{N}_t^{(1)} = 1$ for all detected jumps $\hat{J}_t \neq 0$ and compute $\tilde{\chi}_0^{(1)}$, $\tilde{\rho}^{(1)}$ and hence $\tilde{\xi}^{(1)}$ by solving simultaneously the unconditional mean and autocorrelation noted above. Utilize these coefficients to construct a candidate time series of $\tilde{\chi}_t^{(1)}$. Then impose these model coefficients with the computed $\lambda_z^{(J)}$, $\tilde{\mu}_j^{(1)}$, and $\tilde{\sigma}_j^{(1)}$ on the following log-likelihood:

$$\begin{aligned} \mathcal{L}(\tilde{R}_t; N_t = j, \Omega_{t-1}) &= \frac{1}{2} \ln(2\pi(\hat{h}_t^{(J)} + j(\tilde{\sigma}_j^{(1)})^2)) \\ &\quad + \frac{-(R_t - \tilde{\lambda}_z^{(J)} \hat{h}_t^{(J)} - \lambda_y \tilde{\chi}_t^{(1)} - \tilde{\chi}_t^{(1)}(j + \tilde{\mu}_j^{(1)})^2)^2}{2(\hat{h}_t^{(J)} + \tilde{\sigma}_j^{(1)})^2}. \end{aligned} \quad (3)$$

The likelihood function has only one parameter: the jump compensator λ_z . Therefore,

$$(\hat{\lambda}_z^{(1)}) := \arg \max_{v, \lambda_z} \sum_{t=1}^T \mathcal{L}(\tilde{R}_t; N_t^{(1)} = j, \Omega_{t-1}).$$

Step 3: Compute the likelihood score for each observation $\mathcal{L}(\tilde{R}_t; \hat{\theta}^{(1)})$ at the optimum, where $\hat{\theta}^{(1)} = \{\lambda_z^{(J)}, \tilde{\mu}_j^{(1)}, \tilde{\sigma}_j^{(1)}, \hat{\lambda}_z^{(1)}\}$ is the collection of parameters, the superscript in brackets defining those parameters that will be updated in the following steps. Define an arbitrary threshold $\ln(c_a)$, where c_a is a lower bound, e.g., 1% or 0.1%. For these observations add +1 to the jump counter, $\tilde{N}_t^{(2)} = \tilde{N}_t^{(1)} + 1$ if $\mathcal{L}(\tilde{R}_t; \hat{\theta}^{(1)}) < \ln(c_a)$. For days with multiple jumps we assume that the jumps are of equal size. We therefore compute a new mean and standard deviations for the jumps denoted $\tilde{\mu}_j^{(2)}$ and $\tilde{\sigma}_j^{(2)}$ and subsequently $\tilde{\chi}_0^{(2)}$, $\tilde{\rho}^{(2)}$ and hence $\tilde{\xi}^{(2)}$ from the time new series $\tilde{N}_t^{(2)}$.

Step 4: Maximize the new updated log-likelihood function $\sum_{t=1}^T \mathcal{L}(\tilde{R}_t; N_t^{(2)} = j, \Omega_{t-1})$ if it is greater than $\sum_{t=1}^T \mathcal{L}(\tilde{R}_t; N_t^{(1)} = j, \Omega_{t-1})$ at the optimum, then repeat Steps 2 and 3. If it is equal or lower, discard the new parameters $\hat{\theta}^{(2)}$ and keep $\hat{\theta}^{(1)}$.

In the Online Supplement, we provide a proof, under a fairly general data-generating process and some Monte Carlo studies, to demonstrate the consistency of the iterative scheme above. Notice that the penalty for over-guessing $\tilde{N}_t^{(i)}$ for any given day spills over in subsequent days through the updated $\tilde{\chi}_t^{(i)}$. This process is mediated by the autoregressive parameter ρ . We set c_a to be 0.1% for our models. For example, a time t , log-likelihood of below -6.908 triggers an additional candidate jump. Additionally, notice that this approach only works well when the true average jump size is less than the true standard deviation $\mu_j < \sigma_j$. Thus, small jumps are not generating log-likelihoods greater than one standard deviation away from μ_j and erroneously being awarded extra jumps and hence degrading the total log-likelihood. However, our results indicate that the mean detected jumps are between -0.12% and $+0.7\%$ with standard deviations of between 1.5% and 9% —that is, over one order of magnitude larger.

The Chosen GARCH Models

While jump models have no closed-form solution for the option price, the risk-neutral dynamics for certain types of discrete-time models can be explicitly determined. For GARCH models there is a significant restriction on the autoregressive form of the volatility model to ensure that (a) the risk-neutral dynamics can be easily derived and (b) the resultant option price has a closed-form solution.

This natural limitation is discussed in Christoffersen et al. (2008), who compare a persistent GARCH model with an AR-J model with a more standard GARCH specification with mean reversion. We implement this experimental design. Our major contribution is to propose alternative estimation procedures over the domain of models and provide evidence from a far larger pool of equities and equity indexes for this analysis.

Let $\tilde{h}_t = \int_{\tau_0}^{\tau_t} \sigma_s^2 ds$ be the daily quadratic variation. We assume that \tilde{h}_t is a discrete GARCH process varying

at a daily timescale, and we follow Christoffersen et al. (2008) by setting the spot variance process for the GARCH and CGARCH to be, respectively:

$$h_{t+1} = \omega + \beta h_t + \alpha (z_t - \gamma h_t^{1/2})^2, \quad (4)$$

and

$$h_{t+1} = q_t + \beta(h_t - q_t) + \alpha((z_t - \gamma_1 h_t^{1/2})^2 - (1 + \gamma_1^2 q_t)), \quad (5)$$

$$q_t = \omega + \rho q_t + \varphi((z_t^2 - 1) - 2\gamma_2 h_t^{1/2} z_t). \quad (6)$$

The noise term z_t , in our notation, in the absence of jumps, will be $\int_{\tau_0}^{\tau_t} W(s) ds \equiv W(\tau_t) - W(\tau_0)$ and $r_t = s(\tau_t) - s(\tau_0)$. Therefore, over a single day, the continuous part of the price is a geometric Brownian motion \mathcal{GBM} with constant volatility. For option pricing purposes, the drift term, in the absence of jumps, is

$$\tilde{R}_{t+1} = r + \lambda h_{t+1} + h_{t+1}^{1/2} z_{t+1}, \quad (7)$$

where r is the daily return on a risk-free instrument of same maturity as the option and λ is a coefficient that determines the risk premium. The without-jumps models are fitted using a standard maximum likelihood estimator assuming Gaussian innovations. Hence, $\mathcal{L}_{t_0, T}$ is the Gaussian log-likelihood function for a sample t_0, T , with the vector θ collecting the relevant collection of parameters $\theta = \{\lambda, \alpha, \beta, \gamma, \omega\}$ for the GARCH model and $\theta = \{\lambda, \alpha, \beta, \gamma_1, \gamma_2, \omega, \rho, \varphi\}$ for the CGARCH model. Finally, we denote by $\mathcal{L}_{t_0, T}^*$ the evolution of the log-likelihood at its maxima.

The AR-Jump with GARCH Density Function

Similarly to the volatility model, we assume again that the process driving the intensity of these jumps is autoregressive at a daily frequency, but uniform at the intraday timescale. Let jumps J_t be normally distributed with mean and standard deviation μ_j and σ_j respectively. The distribution of jumps is assumed not to be time varying. For a given day, the arrival rate per intraday block of time is denoted $\tilde{\chi}_t$. From our previous notation, there is a process $\tilde{\chi}_t$ which we assume to be $\tilde{\chi}_t \sim \text{AR}(1)$. We can think of $\tilde{\chi}_t$ as the constant rate of jump arrivals over a single day and it is therefore

scalable to the intraday grids that we specified above— for instance $\bar{\chi}_t$ is the five minute arrival rate when Δ is $5/(24 \times 60)$. $\tilde{\chi}_t$ by contrast is fixed to being the daily intensity rate.

The additivity property of Poisson processes implies that $n_t \bar{\chi}_t = \tilde{\chi}_t$, that is, while $\tilde{\chi}_t$ is autoregressive at a daily timescale, $\bar{\chi}_t$ is constant over a day. As such, the probability mass of a jump arrival is continuously and equally distributed over the trading day. Therefore, during a day, the price process $s(\tau)$ will be a jump diffusion process combining a continuous diffusion with constant volatility and a càdlàg process with constant jump arrival. This follows the implicit assumptions of Tauchen and Zhou (2011), who assumed at most one jump a day. More complex models that permit intraday volatility are of course easily specified, but these have two disadvantages. The first is the need for persistence over many days in the volatility and intensity processes. Second, for a model that captures both intraday and inter-day volatility and intensity persistence, the option pricing model becomes far less tractable with a questionable degree of improvement. We therefore assume an autoregressive modeling structure to the evolution of $\tilde{\chi}_t$ which is at a daily frequency, within the day. The likelihood of a jump occurring is the same for any given fraction of that day:

$$\tilde{R}_{t+1} = r + \tilde{\lambda}_z h_t + \tilde{\lambda}_y \tilde{\chi}_t + \sqrt{h_t} z_{t+1} + \gamma_t + \mu_j \tilde{\chi}_t, \quad (8)$$

$$\tilde{\chi}_t = \tilde{\chi}_0 + \rho \tilde{\chi}_{t-1} + \xi \eta_{t-1}, \quad (9)$$

where Ω is the information set for the filtration of the intensity process, whereby $\mathbb{P}(N_t = j; \Omega_{t-1}) = e^{-\tilde{\chi}_t} \tilde{\chi}_t^j / j!$. The disturbance term η_t is determined recursively such that $\eta_t = \mathbb{E}(N_{t-1}; \Omega_{t-1}) - \tilde{\chi}_{t-1}$. Similarly $\mathbb{E}(N_t; \Omega_t) = \sum_{j=0}^{\infty} j \mathbb{P}(N_t = j; \Omega_t)$ and this is derived recursively from $\mathbb{P}(N_t = j; \Omega_t) = \mathbb{L}(\tilde{R}_t; N_t = j, \Omega_t) \mathbb{P}(N_t = j; \Omega_{t-1}) / (\sum_{j=0}^{\infty} \mathbb{L}(\tilde{R}_t; N_t = j, \Omega_{t-1}) \mathbb{P}(N_t = j; \Omega_{t-1}))$. Where the filtration function is given by

$$\begin{aligned} & \mathbb{L}(\tilde{R}_t; N_t = j, \Omega_{t-1}) \\ &= 2\pi(h_t + j\sigma_j^2)^{1/2} e^{-\frac{(\tilde{R}_t - \tilde{\lambda}_z h_t - \lambda_y \tilde{\chi}_t - (\mu_j - v\sigma_j^2)\tilde{\chi}_t)^2}{2(h_t + \sigma_j^2)}} \end{aligned} \quad (10)$$

and γ is a compound Poisson random variable with mean μ_j and variance σ_j^2 . Note that we now have two risk premium coefficients, namely $\tilde{\lambda}_z$ and $\tilde{\lambda}_y$.

Risk-Neutral Dynamics

For the pricing of options we must recast the parameters of the empirical model to provide a set of risk-neutral dynamics. This is equivalent to a change of measure in a constant volatility framework. Given some of the subtle adjustments to our framework, it is worth re-stating the risk-neutral pricing models from Christoffersen et al. (2008). In summary, we compare three basic pricing approaches: the GARCH and the CGARCH, which are computed using the Kolmogorov backward method, and the ARJ-GARCH, for which the price is computed via Monte Carlo simulation. In each case, we write down the risk-neutral price and then proceed to derive the functional form or infer the price of the option. The risk-neutral dynamics for the Heston and Nandi (2000) model are given by

$$\tilde{R}_{t+1} = r + \lambda h_{t+1} + \sqrt{h_{t+1}} z_{t+1}^* \quad (11)$$

$$h_{t+1} = \omega + \beta h_t + \alpha (z_t^* - \gamma^* h_t^{1/2})^2, \quad (12)$$

where $\gamma^* = \gamma + \lambda + 1/2$ and $z_t^* \sim \mathcal{N}(0,1)$. In a noteworthy contribution to this area, Christoffersen et al. (2008) demonstrate that the CGARCH risk-neutral dynamics are given by

$$\tilde{R}_{t+1} = r + \lambda h_{t+1} + \sqrt{h_{t+1}} z_{t+1}^* \quad (13)$$

$$h_{t+1} = q_{t+1} + \beta^* (h_t - q_t) + \alpha((z_t^* - \gamma_1^* h_t^{1/2})^2 - (1 + \gamma_1^{*2} h_t)) \quad (14)$$

$$q_{t+1} = w + \rho^* q_t + \varphi((z_t^* - \gamma_2^* h_t^{1/2})^2 - (1 + \gamma_2^{*2} h_t)), \quad (15)$$

where the risk-neutral parameters (denoted by the (*) term) are defined by

$$\beta^* = \beta + \alpha(\gamma_1^{*2} - \gamma_1) + \varphi(\gamma_2^{*2} - \gamma_2) \quad (16)$$

$$\rho^* = \rho + \alpha(\gamma_1^{*2} - \gamma_1) + \varphi(\gamma_2^{*2} - \gamma_2) \quad (17)$$

$$\gamma_i^* = \gamma_i + \lambda + 1/2, \quad i \in \{1, 2\}. \quad (18)$$

The final baseline option pricing model is the ARJ-GARCH. The risk-neutral dynamics for the mean, variance, and jump equations are as follows:

$$\tilde{R}_{t+1} = r + \lambda_z h_{t+1} + \lambda_y \tilde{\chi}_{t+1}^* + \sqrt{h_{t+1}} z_{t+1}^* + \gamma_{t+1}^* - \mu_j \tilde{\chi}_{t+1}^* \quad (19)$$

$$h_{t+1} = \omega + \beta h_t + \alpha(z_t^* + \gamma_t^* h_t^{-1/2} - \gamma^* h_t^{1/2})^2 \quad (20)$$

$$\tilde{\chi}_t^* = \tilde{\chi}_t e^{\frac{1}{2}v^2\sigma_J^2 + v\mu}, \quad (21)$$

where the evolution of $\tilde{\chi}_t$ is from (8), v is the market price of risk and $\mu^* = \mu_J - v\sigma_J^2$, $\gamma^* = \gamma + \frac{1}{2} + \lambda_y^*$ and $\lambda_y^* = (\lambda_y - \mu_J) \exp(-1/2v^2\sigma_J^2 - v\mu_J)$. The noise components z_t^* and γ_t^* are respectively a normally distributed random variable under a risk-neutral measure transformed from z_t and a Poisson jump with normally distributed jump size $\gamma_t^* \sim \mathcal{N}(\mu_J^*, \sigma_J)$ with a Poisson count $N_t^* \sim \text{Poisson}(\tilde{\chi}_t^*)$. It is worth noting that the term $(\lambda_y^* - \mu_J^*)\tilde{\chi}_t^*$ forces the discounted price process to be a martingale.

We now have an unknown parameter v , that needs to be derived from the maximum likelihood estimates of the model specified in equation 8. The option prices for the autoregressive jump models in mean and variance, by construction, must be derived from the risk-neutral prices rather than using a complete market approach. We follow Pan (2002), albeit modified to jumps in the first and second moments and including the explicit autoregressive jump construction. This has the disadvantage that the value of v needs to be calibrated by simulation rather than via a numerical integration approach which is, of course, computationally quicker.

Unfortunately, analytical long-run moments for this type of process are not possible to compute (if they were then the option could be solved without recourse to simulation). We need thus to compute first ‘‘physical’’ maximum likelihood estimates and then impute v . Christoffersen et al. (2008) suggest that v can be ‘‘solved for numerically from the physical MLE estimates’’ and we argue this to be possible when the intensity process χ_t is time invariant. However, experimentation indicates that relying on the approach of Pan (2002) provides far better results.

The Pan (2002) approach is relatively straightforward albeit computationally intensive.

- Step 1: Estimate the observed data-generating process from the estimation window of the underlying sample.
- Step 2: Choose a lower bound on $v^{(1)}$ and impute the risk-neutral parameters $\lambda_y^{(1)*}$, $\mu_J^{(1)*}$, and $\gamma^{(1)*}$. Generate the time series $\chi_t^{(1)*}$ and $h_t^{(1)}$ under

the risk-neutral dynamics. Next collect a daily sample of very near maturity, at-the-money, options.

- Step 3: Pan (2002) recommends less than 25 days. We posit that a fair test is to collect the average price of 15 to 25 (for stocks this is required for coverage) day at the money option prices for each day in the sample and compute the average maturity $T^* - t$. Next generate 10,000 draws for each day starting from $\chi_t^{(1)*}$ and $h_t^{(1)}$ and compute the undercounted estimated price of each option, under $v^{(1)}$. Using the rule:

$$\begin{aligned} \tilde{P}_t^{(1)*} &= \frac{1}{10,000} \sum_{i=1}^{10,000} \max(S_{t,T}^{(1)*} - K, 0) \\ &\rightarrow \int_0^\infty \max(S_T^{(1)*} - K, 0) d\mathbb{Q}^{(1)}_s \end{aligned}$$

and hence the option price from $C_t^{(1)} = e^{-r(T^*-t)}$, where $\mathbb{Q}^{(1)}$ is the risk-neutral measure under $v^{(1)}$ and collect the root mean squared log difference between the observed C_t and $C_t^{(1)}$. Recall that we denoted this as $\mathcal{M}^{(1)}$.

- Step 4: Compute a new market price of jump risk from the rule $v^{(a+1)} = v^{(a)} + \mathcal{E}$, for the a iteration. Where \mathcal{E} is a computationally manageable step size and repeat Step 2.

We suggest two termination conditions: first iterative, if $\mathcal{M}^{(a+1)} > \mathcal{M}^{(a)}$, then record $v = v^{(a)}$. However, as noted in Pan (2002), the error in the first moment can be problematic. We therefore advocate searching over a range v^{\min} to v^{\max} , which is easily encoded in parallel, and then identifying the $v^{(a)}$ smallest pricing error $\min(\mathcal{M}^{(a)})$.

Pricing the Options Using CGARCH and GARCH Models

As discussed above, for the GARCH and CGARCH models, closed-form solutions exist for the pricing of the options. The pricing formulas are in the form of identification of the functions $N_S(\vartheta, \theta)$ and $N_K(\vartheta, \theta)$, where θ is the vector of parameters of the GARCH ($\theta = \{\lambda, \alpha, \beta, \gamma, \omega\}$) and CGARCH ($\theta = \{\lambda, \alpha, \beta, \gamma_1, \gamma_2, \phi, \rho, \omega\}$) models and ϑ is the vector of parameters of the option

$\vartheta = \{t, T, K, S, r\}$. In both cases, the closed-form solution is derived using the characteristic function. In particular, for each case the local characteristic function $f^*(t, T, \phi)$ for each option is derived using Kolmogorov's backward method. The price of the call option is of the form

$$C = e^{-r(T-t)} \mathbb{E}_t^*(\max(S_T - K, 0)) \\ = S_t \mathbf{N}_S(\vartheta, \theta) - Ke^{-r(T-t)} \mathbf{N}_K(\vartheta, \theta),$$

where \mathbb{E}_t^* is the time t conditional expectation under risk-neutral distribution. For the GARCH model, Heston and Nandi (2000) derive the moment generating function for any number of lags in the AR and MA components of the model. The functions $\mathbf{N}_S(\vartheta, \theta)$ and $\mathbf{N}_K(\vartheta, \theta)$ are defined as follows:

$$\mathbf{N}_S(\vartheta, \theta) = \frac{1}{2} + \frac{1}{\pi} \int_0^\infty d_S(\phi) d\phi, \\ \mathbf{N}_K(\vartheta, \theta) = \frac{1}{2} + \frac{1}{\pi} \int_0^\infty d_K(\phi) d\phi, \quad (22)$$

where

$$d_S(\phi) = \operatorname{Re} \frac{K^{-i\phi} f^*(t, T, \phi + 1)}{i\phi S_t e^{r(T-t)}}, \quad \text{and} \\ d_K(\phi) = \operatorname{Re} \frac{K^{-i\phi} f^*(t, T, \phi)}{i\phi S_t e^{r(T-t)}}, \quad (23)$$

where $f^*(t, T\phi) = f(t, T, i\phi)$, $f(\phi)$ is the moment-generating function, over ϕ moments for the candidate process. Let ϕ be the unconditional moments; for the first-order version of the GARCH model, the moment-generating function is given by $f(t, T, \phi) = S_t^\phi e^{A_t + B_t h_{t+1}}$, where

$$A_t = A_{t+1} + r\phi - B_{t+1}\omega - \frac{1}{2} \ln(1 - 2\alpha B_{t+1}), \\ B_t = \phi(\lambda + \gamma) - \frac{1}{2} \gamma^2 + \beta B_{t+1} \quad (24)$$

with the time T boundary condition is that $A_T = B_T = 0$. Heston and Nandi (2000) provide extensive coverage of the GARCH approach utilizing the Kolmogorov's backward induction equation to identify the above moment-generating function. Christoffersen et al. (2008) extend the approach for the GARCH to the CGARCH form. Setting $f(t, T, \phi) = S_t^\phi e^{A_t + B_{1,t}(h_{t+1} - q_{t+1}) + B_{2,t}q_{t+1}}$ to be the

moment-generating function with recursive coefficients, then

$$A_t = A_{t+1} + r\phi - (\alpha B_{1,t+1} + \phi B_{2,t+1}) \\ - \frac{1}{2} \ln(1 - 2\alpha B_{1,t+1} - 2\phi B_{2,t+1}) + \omega B_{2,t+1} \quad (25)$$

$$B_{1,t} = \beta B_{1,t+1} + \lambda\phi + 2 \frac{(\alpha\gamma_1 B_{1,t+1} + \phi\gamma_2)^2}{1 - 2\alpha B_{1,t+1} - 2\phi B_{2,t+1}} \quad (26)$$

$$B_{2,t} = B_{2,t+1} + \lambda\phi + 2 \frac{(\alpha\gamma_1 B_{1,t+1} + \phi\gamma_2)^2}{1 - 2\alpha B_{1,t+1} - 2\phi B_{2,t+1}} \quad (27)$$

and terminal conditions $A_T = B_1$, $T = B_2$, $T = 0$, effectively adding the component structure to the original GARCH specification. However, the GARCH model has a deliberate structure that imposes certain regularity conditions on the characteristic function (see the technical Appendix in Heston and Nandi 2000), and from the complexity of the CGARCH component equation it is effectively impossible to recover analytically confirmation that the CGARCH will converge in an identical fashion. In order to address this, we will explore the CGARCH pricing function in more detail using numerical examples.

Numerical Stability of the Pricing Algorithms

For the GARCH and CGARCH models the objective is to compute the functions $\mathbf{N}_i(\vartheta, \theta)$, $i \in \{S, K\}$ via numerical quadrature. The option pricing literature that encompasses the popular Heston (1993) approach proffers a considerable discussion on the numerical stability of such an approach; see Albrecher et al. (2006) for a discussion of the "Little Heston Trap" problem that can substantiate itself in the computation of options under stochastic volatility. Notice that the major issue is the branch cut in the complex plane for the solutions to the characteristic function. The Heston model is arguably the simplest nontrivial model that can be computed via fast Fourier transform or numerical quadrature of the characteristic function, and the existence of two roots in the complex plane for this model poses several problems for numerical implementation. The Heston and Nandi (2000) GARCH model essentially collapses to the Heston (1993) model as $\Delta t \rightarrow 0$ and similarly

provides a well-defined branch cut that we can identify. However, the CGARCH model is more complicated in this sense and numerical implementation poses problems. The problem of the branch cut in the evaluation of the function is that at the cut numerical evaluations can exhibit substantial errors, an obvious issue being singularities that result in indefinite evaluations.

To overcome this issue, some publicly available algorithms for the CGARCH pricing model appear to solve this by arbitrary truncation at $\phi = 10$. However, this truncation also causes problems that deteriorate the accuracy of the option price, as for many out-of-the-money strikes the function has not yet fully converged to zero at $\phi = 10$. Several quadrature methods are commonly used to solve the types of integrals found in equation 22, and standard implementation techniques are available in most numerical analytical tools such as MatLab or Mathematica. To illustrate that the effect stems from the moment-generating function and not the quadrature method we implement three procedures: Simpson's rule, the Gauss-Konrod rule, and the trapezoidal or brute force integration. We also have implemented the option pricing method in both Mathematica (which uses variable or arbitrary precision arithmetic) and MatLab that uses IEEE standard double precision arithmetic.

From a practical standpoint, using an arbitrary precision arithmetic to price a large number of options is effectively impossible, as each option requires a substantial amount of computation time to derive the price, even using a large cluster. Hence, we demonstrate the instability of the CGARCH model for a pair of options. In Exhibit 5 we present CGARCH parameter estimates for our models compared with those found in see the technical Appendix Heston and Nandi (2000) and Christoffersen et al. (2008) for the S&P 500 index (highlighted in bold red for our estimated and bold black for Heston and Nandi (2000) and Christoffersen et al. (2008)). For comparison we will use the parameters from the non-jump-filtered returns. Using these parameters, we analyze the following pair of options: a 100-day at-the-money call and a 300-day in-the-money call where $S/K = 1.3$. We will assume that in all cases the conditional variance is $h_t = 0.35^2$ and for the CGARCH the component volatility is $q_t = 0.0179^2$, the average q_t from our CGARCH model with no jump filtration. Let

$$\begin{aligned}\mathcal{I}_S(\bar{\phi}) &= \frac{1}{2} + \frac{1}{\pi} \int_0^{\bar{\phi}} d_s(\phi) d\phi, \\ \mathcal{I}_K(\bar{\phi}) &= \frac{1}{2} + \frac{1}{\pi} \int_0^{\bar{\phi}} d_K(\phi) d\phi.\end{aligned}\tag{28}$$

In Exhibit 1 we observe, over a range of initial stock-prices, the explosive point for an example set of evaluations of $d_s(\phi)$. While the spectral functions $d_s(\phi)$ and $d_K(\phi)$ for the GARCH model converge to zero after a finite number of iterations, the CGARCH versions of the functions $d_s(\phi)$ and $d_K(\phi)$ do not. In fact, they stabilize and then after a finite number of iterations the function destabilizes and explodes. A further complication is that for the two different option models the point at which the CGARCH functions $d_s(\phi)$ and $d_K(\phi)$ destabilize diverges substantially, and consistently deriving this threshold has proved elusive. Our approach is to systematically vary the upper limit of integration and then use a Haar wavelet (a local average) to detect the point of relative flatness and hence the correct values of $\mathbf{N}_{i \in \{S, K\}}$.


Note that the stability of the estimates deteriorates when jumps are incorrectly omitted from the option pricing model. Consider a data-generating process for $\tilde{R}_t = R_t + J_t$, where R_t is the continuous diffusion component with GARCH volatilities and J_t is the jump component. While the option is biased if the option price is based solely on R_t rather than \tilde{R}_t , the contamination of the auto-regressive jumps causes greater problems for the CGARCH model in correctly pricing options, and this contamination has differential effects across maturities and money-ness. We will show that the jump-filtered GARCH and CGARCH models perform at least as well as, and in many cases better than, GARCH and CGARCH models estimated on unfiltered returns. The bias in the price also will vary with the properties of the jump model. We find that the deterioration in the performance of the CGARCH when jumps are omitted is substantial.

Out-of-Sample Forecasting

The final aspect of our analysis is an ‘‘acid-test’’ of the volatility model in comparing the computed call option prices to traded options out-of-sample. We therefore estimate the various GARCH models over

EXHIBIT 1

Numerical Stability of CGARCH Option Pricing Function. Evaluation of $\int_0^{\bar{\phi}} d_s(\phi)d\phi$ with Respect to an Arbitrary Truncation $\bar{\phi}$ the Ordinate Values on the Left Plot are Multiplied by a Factor of 10 for Visualization Purposes



Notes: The left-hand plot illustrates the evaluations of $\int_0^{\bar{\phi}} d_s(\phi)d\phi$ with respect to the upper bound $\bar{\phi}$. Many implementations of these types of model truncate $\bar{\phi}$ to around the 10th moment. We can see from this simple example where $\lambda = -0.2049$, $\beta = 0.9132$, $\alpha = 1.3201e-07$, $\gamma_1 = 415$, $\omega = 6.8328e-07$, $\rho = 0.9855$, $\phi = 2.8006e-06$, $\gamma_2 = 2.0134e-10$, $h_t = 0.35^2$ and $q_t = 0.0179^2$. The option being priced is 50 days to maturity and the risk-free rate is set to 0.25%. Note that the ordinate values for the right-hand plot are already in the range $\pm 1e199$ when $\bar{\phi} = 247.5$. This provides a fairly complicated numerical problem to solve, as the function slowly deviates toward the singularity, an error that has been noted in this type of approach to pricing options.

the sample period up to the onset of our options data sample in 2010. An important aspect of pricing options is to ensure that intraday effects do not add noise and/or bias to the results. Our options are sampled each day from the OPRA data feed at 4 p.m. and matched to the underlying assets price at that time. This is considered to be time t . The conditional variances and jump intensities used in the option pricing are taken from the previous day's closing stock or index return $t - 1$, and are therefore ex ante and not ex post. Hence, we assume that h_t and χ_t are known and that the parameters of the model rolling forward are taken from the preceding estimation window, from 1996 to 2010, which is the full history of the available tick data from Thomson-Reuters. To test for parameter instability we implement the fixed window, rolling window, and recursive tests described in Giacomini

and Rossi (2010) and find very little variation for most of the models except the CGARCH when estimated on the raw returns, not filtered for jumps. Here the rolling window reveals a substantial change in the ω coefficient in the q_t equation during 2010 for the S&P 500 index, and several stocks appear to violate the bootstrapped maximal fluctuation test from Giacomini and Rossi (2010) critical boundary at 90%, but not at 95%. Therefore, we still adopt the fixed estimation window coefficients for the CGARCH, to maintain parity with the other models.

Call Option Loss Functions and Experiments

Let $C_t(\vartheta_t)$ and $\tilde{C}_t(\vartheta_t, \theta_{t-1})$ be the actual recorded call option price and the model predicted call option price for day t . The parameter vectors are partitioned

as follows. First, the option-specific information is given by $\vartheta_t = \{S_t, K, r, T\}$ where S_t is the actual stock price at the time of trade, K is the strike price, r is the current risk free rate, and T is the maturity. The one-day-lagged information θ_{t-1} is model specific. For the GARCH model this is $\theta_{t-1} = \{\omega, \beta, \alpha, \gamma, h_t\}$, for the CGARCH model $\theta_{t-1} = \{\omega, \rho, \alpha, \beta, \gamma_1, \gamma_2, h_t, q_t\}$ and for the ARJ-GARCH $\theta_{t-1} = \{\omega, \alpha, \beta, \gamma, \lambda_\gamma, \nu, \sigma_j, \mu_j, h_t, \tilde{\lambda}_t\}$. Following Christoffersen and Jacobs (2004), we set up a loss function approach by setting $F_j(C_{j,t}(\vartheta_t), \tilde{C}_{j,t}(\vartheta_t, \theta_{t-1})) = F(\delta_{j,t})$ where j indexes the $M(\vartheta)$ call options with $\vartheta \in \Theta$ characteristics in our sample. The summary loss statistics are the root mean and root median forecast errors:

$$\mathcal{F}_\mu(\vartheta) = \sqrt{1/M(\vartheta) \sum_{j=1}^{M(\vartheta)} F(\delta_{j,t})} \quad (29)$$

$$\mathcal{F}_m(\vartheta) = \sqrt{\text{median}(F(\delta_{j,t}))}. \quad (30)$$

We choose the loss function $F(\delta_{j,t})$ to be in terms of square of the log differences:

$$F(\delta_{j,t}) = (\ln C_{j,t}(\vartheta_t) - \ln \tilde{C}_{j,t}(\vartheta_t, \theta_{t-1}))^2 \quad (31)$$

We then assign each option to a moneyness-maturity bin and compute the root mean of the quadratic losses.

There are of course a very large number of combinations that can be studied within an empirical analysis of this form. In order to present the results in a relatively contained manner we choose 20 moneyness-maturity bins with four maturities, $0 < T - t \leq 50$, $50 < T \leq 100$, $100 < T - t \leq 500$, $T > 500$, and five stratifications of moneyness, $0 < S/K \leq 0.75$, $0.75 < S/K \leq 1$, $1 < S/K \leq 1.25$, $1.25 < S/K \leq 1.5$, and $S/K > 1.5$.

We restrict ourselves to presenting six models. We also implement a further pair of models: the generalized affine realized volatility model of Christoffersen et al. (2012) and the doubly stochastic poisson model of Scott (1997), calibrated only to historical options data. However, the results for both models were very markedly inferior to the CGARCH model estimated using unfiltered returns, the worst-performing model from our chosen group. Code and results for these models are available in our online supplementary material.

- #1 GARCH Estimated by Maximum Likelihood on jump-filtered returns. Option prices via numerical integration of $d_{t \in \{S,K\}}$.
- #2 CGARCH Estimated by Maximum Likelihood on jump-filtered returns. Option prices via numerical integration of $d_{t \in \{S,K\}}$.
- #3 GARCH with AR-Jumps Estimated via the iterative approach from mixed frequency data as outlined above.
- #4 GARCH Estimated by Maximum Likelihood on unfiltered returns. Option prices via numerical integration of $d_{t \in \{S,K\}}$.
- #5 CGARCH Estimated by Maximum Likelihood on unfiltered returns. Option prices via numerical integration of $d_{t \in \{S,K\}}$.
- #6 GARCH with AR-Jumps Estimated by Maximum Likelihood on unfiltered returns. Options priced via Monte Carlo simulation.

The first three models mirror essentially Christoffersen et al. (2008) for both stocks and the S&P 500 index. We disaggregate the presentation of the results for the S&P 500 index options from the individual stocks. Furthermore, to illustrate the consistency across individual stocks, we present two cases: IBM and Johnson and Johnson (Code JNJ). We then report the mean squared errors for all the individual stocks combined.

DATA AND PREPROCESSING

When conducting a study of this type the probity and sources of the data are of critical importance. It is therefore instructive to carefully detail the data collection approach to illustrate any potential sources of bias. We collect our stock and index data from the Thomson-Reuters Tick History (TRTH) data service, which in turn pulls data directly from the exchange feeds. For the pre-2005 period this is by exchange and post-2005 this is from the National Market System (NMS) for both the underlying stocks and the options via OPRA.

When conducting intraday volatility studies, the econometrician has the discretion of using the mid-price halfway between the best bids and best asks (the inside spread) or the actual traded prices. We choose the actual

trades, as (a) this is the price that the option will be priced against and (b) this type of tick data is consistent with the stock index data. Note that by license we are not permitted to release source tick data. Nevertheless, the time series of daily returns, realized volatilities, bipower variation, tripower quartilities, and inferred jumps for our one-minute data is available in the Online Supplement.

The historical stock and index sample pulls data available from January 1, 1996, and ends on March 6, 2014, the entirety of the TRTH available history, covering 17 years for 20 individual stocks sampled from the components of the Dow Jones Industrial Average and the intraday S&P 500 index. Notably, some stocks (for instance Chevron, Verizon, and Home Depot) start slightly later in the data feed. The estimation of the underlying models is conducted on the daily aggregate returns, denoted \tilde{R}_t (see above for a detailed explanation).

The options data are sampled from the OPRA data feed via TRTH from January 1, 2010, to March 6, 2014. It is important to point out that the source options data from the OPRA NMS is vast. Therefore, our filtering approach is as follows. We pull only the options traded in the five minutes after 4 p.m. EST weekdays. Note that the vast majority of options trades are conducted in Chicago on the CME, while the underlying tickers are mostly traded in New York and New Jersey, so the time zones are different. We time stamp all the trading data relative to GMT to ensure that stocks and options are synchronized. All live maturity dates from January 2, 2010, to March 6, 2014, are included. For each available contract type (each code contains a maturity date and a strike price) for our underlying assets we compute the open-interest weighted price of all trades within the five-minute time period of interest. This is accomplished by an algorithm within the data service. Hence, the number of source contracts used to compute open-interest-weighted price is unknown. However, pulling the raw data for the S&P 500 index options for a few days indicated that it is in the thousands for the at-the-money options and in the hundreds for the out-of-the-money options. Therefore, while for brevity we refer to the data as 347,612 contracts, it is more appropriate to refer to it as 347,612 maturity-strike-ticker-code-day averages for the five-minute sample in question. All options are restricted to the European type from the data feed. Exhibit 2 provides details of the sample characteristics.

Despite the active nature of the equity options market, some stocks have missing or very few maturity-

moneyness combinations; this is illustrated in Exhibit 3, which cuts up the sample by contract population in bins. It should be noted that the S&P 500 index options have a much wider range of available strike prices registered on the OPRA-NMS system. Accordingly, over one-third of the sample comes from this contract type. This source of data is roughly that which provides the end-of-day prices reported in the OptionMetrics data set. However, we choose 4 p.m. to eliminate potential end-of-day effects. This appears to be consistent with market practice, most notably for individual stocks.

The historical data set, therefore, consists of over 2 billion trades and/or updates from the S&P 500 index and 20 stocks randomly selected from the Dow Jones 30. The GARCH models are fitted on the filtered data up to January 1, 2010. However, for the remaining data we compute the forward recursion of the models until the end of the sample on March 6, 2014. This is an average of 538,322 trades or updates per day and 93,030 firm days in the sample, making this one of the largest scale studies of its type, on RV alone. The option data set spans January 1, 2010, to March 6, 2014, and consists of 347,612 traded option contracts. Again, to our knowledge this is one of the largest data sets ever constructed for option pricing problems involving Monte Carlo estimations.

ANALYSIS AND ECONOMIC IMPLICATIONS

In this section, we first summarize the parameter estimates in-sample for our new approach to fitting the models to the underlying data using mixed frequencies for the estimation window. Finally, we review the mean and median results for the models out-of-sample for our options data.

In-Sample Model Estimates

We begin our discussion with the maximum likelihood estimations to document the consistency of the estimates across our individual stocks and the index and to demonstrate that our results for the stocks are not driven by one rogue miscalibrated MLE. In Exhibit 3 we provide a summary of our option contracts across firms and the S&P 500 index. The final column, in bold, shows the summary for the total number of contracts for just the firms, and the middle column highlighted in bold presents the number of contracts for the S&P 500 index.

EXHIBIT 2

Sample Characteristics

Codes	Name	Average Ticks per Day	Number of Days in Sample	Total Ticks	No. of Options Contracts	Market Price of Jump Risk (ν in Basis Points)
T	AT&T	35,758	4,641	165,952,878	11,206	209.61
AA	Alcoa	27,545	4,658	128,304,610	10,299	449.09
AXP	American Express	21,061	4,604	96,964,844	10,266	252.64
BA	Boeing	14,371	4,628	66,508,988	12,773	252.63
CAT	Caterpillar	18,470	4,640	85,700,800	8,720	521.46
CVX	Chevron	33,048	3,212	106,150,176	8,948	695.00
KO	Coca Cola	20,694	4,608	95,357,952	8,366	830.60
DIS	Dupont	21,549	4,606	99,254,694	7,899	245.22
DD	Disney	15,097	4,584	69,204,648	9,938	130.17
GE	General Electric	58,028	4,716	273,660,048	7,910	287.85
HPQ	Home Depot	46,298	3,084	142,783,032	8,646	398.66
HD	Hewlett Packard	26,473	4,631	122,596,463	8,473	913.71
IBM	Int. Bus. Mach.	17,168	4,637	79,608,016	22,641	521.31
JNJ	Johnson and Johnson	23,857	4,619	110,195,483	10,524	27.99
MCD	McDonalds	16,187	4,611	74,638,257	7,401	530.29
PG	Procter and Gamble	22,714	4,596	104,393,544	9,401	154.50
PFE	Pfizer	43,148	4,689	202,320,972	8,744	781.19
SPX	S&P 500 Index	2,142	4,543	9,731,106	144,025	46.63
UTX	United Technologies	12,346	4,563	56,334,798	9,604	925.62
VZ	Verizon	34,030	3,526	119,989,780	14,792	178.68
WMT	Walmart	28,338	4,634	131,318,292	7,036	604.56
Total		538,322	93,030	2,340,969,381	347,612	

Notes: For computational reasons, we restrict ourselves to 20 stocks randomly selected from the Dow 30 and the S&P 500 index; note that we use the intraday index, which is updated roughly every 15 seconds during the trading day. Thus, the number of informed updates is relatively stable at just over 2,000. Our data come directly from the Reuters-America feed that computes the update and this appears to contain more than the 1,560 updates expected for the 390 minutes of the trading day. The S&P 500 depository receipt (SPDR), an S&P 500 exchange traded fund, has a higher update frequency. However, this is an "equal-weight" version of the S&P 500 and it is not the reference index for the S&P 500 index options. The codes in column one are used throughout the paper. Average ticks per day are the average number of trades in a day. The number of trading days for each sample includes all days where there are more than 60 price changes. In practice, we have only dropped one day, for AA, in the whole sample, due to the lack of variation in ticks. Total ticks are provided to give an overview of the whole historical data set. Number of options contracts shows the total sample size for 2010 to 2014, for our out-of-sample testing period. Market price of jump risk is computed from the autoregressive jump model maximum likelihood parameters using end-of-day short-maturity near-the-money options, in our estimation window.

We are relatively fortunate in that the number of available contracts for the individual stocks is relatively even across our bins. However, certain bins do stand out. For instance, we can observe a very large number of long maturity options with maturity in excess of 250 days. This must be noted as a critical issue for the S&P 500. In fact, the standardized option strikes mean that the number of available contracts that are out-of-the-money by more than 25% are in short supply. This is a result very similar to that found in Pan (2002) where the available strikes are dependent on the exchanges issuing new standardized strike-maturity codes. For options of longer maturity, more strike-maturity combinations are

issued and hence we have a larger sample for the cross-section of strikes. It is also noteworthy that the bin sizes are not identical. Hence, we can capture more interesting pricing errors for near-the-money and heavily in- and out-of-the-money conditions.

In Exhibit 4 we present the maximum likelihood estimates for the stocks and the S&P 500 index for the jump-filtered data generated from the tick-by-tick returns for the GARCH model. The term $\mathcal{L}_{i,T}^*$ denotes the evaluation of the maximized log-likelihood, when the models are estimated using the assumption of conditional normality. The last column in bold presents for comparison the estimates of Christoffersen et al. (2008)

EXHIBIT 3 Option Contract Data

	AA	AXP	BA	CAT	CVX	DD	DIS	GE	HD	HPQ	IBM	JNJ	KO	MCD	PFE	PG	SPX	T	UTX	VZ	WMT	Stocks
0 < T - t ≤ 75																						
0 < $\zeta \leq 0.75$	238	7	247	1	362	3	2	71	105	150	557	3	14	100	17	3	38	138	146	177	27	2,368
0.75 < $\zeta \leq 0.975$	244	327	465	336	454	243	315	504	436	601	758	326	644	345	581	319	2,554	642	265	412	413	8,630
0.975 < $\zeta \leq 1.025$	36	91	86	120	74	67	101	91	89	89	124	120	134	126	104	107	1,212	103	39	154	100	1,955
1.025 < $\zeta \leq 1.25$	132	276	323	353	275	203	331	299	299	359	429	361	463	319	375	280	4,101	408	149	570	337	6,541
1.25 < $\zeta \leq 1.50$	108	133	229	40	204	167	177	243	155	265	330	154	121	155	213	114	3,058	97	111	428	105	3,549
$\zeta > 1.5$	368	108	270	4	169	111	253	241	68	44	396	213	17	141	322	66	2,312	84	142	220	116	3,353
75 < T - t ≤ 100																						
0 < $\zeta \leq 0.75$	69	0	57	0	118	0	0	45	26	69	153	12	5	50	12	12	21	82	51	111	12	884
0.75 < $\zeta \leq 0.975$	94	90	180	120	181	64	75	375	187	250	245	119	273	146	411	135	1,223	466	121	291	160	3,983
0.975 < $\zeta \leq 1.025$	14	33	43	50	35	26	41	71	43	40	49	48	56	61	81	44	654	84	18	119	44	1,000
1.025 < $\zeta \leq 1.25$	51	103	125	155	108	70	121	230	124	151	151	129	181	112	277	116	1,787	306	85	425	124	3,144
1.25 < $\zeta \leq 1.50$	44	76	66	57	67	59	75	155	57	101	110	58	114	63	153	42	1,232	133	51	318	50	1,849
$\zeta > 1.5$	149	72	102	68	78	70	89	125	60	138	130	79	55	63	120	63	1,121	68	82	186	59	1,856
100 < T - t ≤ 150																						
0 < $\zeta \leq 0.75$	240	0	43	1	255	0	0	50	8	97	715	0	2	26	0	0	64	121	120	97	0	1,775
0.75 < $\zeta \leq 0.975$	250	217	301	229	340	185	174	535	231	314	828	202	354	238	547	224	1,207	666	250	408	250	6,743
0.975 < $\zeta \leq 1.025$	32	77	50	62	58	58	67	93	59	48	135	82	72	90	114	74	374	125	40	158	72	1,566
1.025 < $\zeta \leq 1.25$	144	266	174	212	189	199	234	326	195	181	463	286	248	239	392	249	1,362	427	177	546	226	5,373
1.25 < $\zeta \leq 1.50$	105	177	126	86	143	145	169	213	84	142	361	105	194	127	238	112	1,053	229	108	372	112	3,348
$\zeta > 1.5$	497	165	242	130	179	160	197	191	107	265	397	162	138	64	282	142	2,243	178	169	402	88	4,155
150 < T - t ≤ 250																						
0 < $\zeta \leq 0.75$	1,139	7	342	38	335	12	1	90	54	355	1,744	0	7	0	1	0	148	175	379	93	0	4,772
0.75 < $\zeta \leq 0.975$	704	750	794	589	601	655	721	624	641	819	2,148	592	936	338	644	566	2,564	940	667	727	316	14,772
0.975 < $\zeta \leq 1.025$	105	202	133	174	97	157	174	117	153	127	337	212	194	157	137	194	633	149	111	280	129	3,339
1.025 < $\zeta \leq 1.25$	393	629	457	493	344	515	574	406	468	475	1,217	612	661	395	484	538	2,318	570	423	1,056	381	11,091
1.25 < $\zeta \leq 1.50$	301	291	344	193	257	384	322	293	162	361	931	267	386	182	324	246	1,696	411	289	772	110	6,826
$\zeta > 1.5$	1,256	246	523	186	364	181	316	259	188	491	1,089	287	181	10	432	284	5,064	356	432	955	104	8,140
T - t > 250																						
0 < $\zeta \leq 0.75$	1,242	396	866	765	413	168	334	216	538	201	1,681	301	312	555	216	378	5,388	648	251	679	488	10,648
0.75 < $\zeta \leq 0.975$	650	1,669	1,610	1,454	1,166	1,246	1,410	702	1,333	573	2,358	1,647	858	1,004	741	1,711	21,822	1,344	1,295	1,664	1,132	25,567
0.975 < $\zeta \leq 1.025$	47	412	273	357	192	314	352	189	311	102	459	428	165	310	198	390	5,440	238	227	291	260	5,515
1.025 < $\zeta \leq 1.25$	356	1,391	945	1,144	680	1,032	1,195	426	1,074	380	1,654	1,509	582	941	499	1,315	19,239	846	1,106	1,159	776	19,010
1.25 < $\zeta \leq 1.50$	178	906	937	518	508	588	756	361	643	331	1,239	854	407	477	297	636	14,394	528	894	775	364	12,197
$\zeta > 1.5$	1,113	1,149	2,420	785	702	817	1,362	369	748	954	1,453	1,356	592	567	532	1,041	39,703	644	1,406	947	681	19,638
Total	10,299	10,266	12,773	8,720	8,948	7,899	9,938	7,910	8,646	8,473	22,641	10,524	8,366	7,401	8,744	9,401	144,025	11,206	9,604	14,792	7,036	203,587

Notes: A summary of the number of options contracts for each of the 25 domains moneyness and maturity "bins" used in our sample. It is important to note that these are not individual trades, but means of prices taken over a 10-minute period before 3 p.m. Therefore, each number represents the number of contracts available within each bin. The mean values of the individual trades are computed by CME. A zero indicates that no contract applicable to the individual bin is available. A color version of this exhibit appears in the online edition of this article.

EXHIBIT 4

Maximum Likelihood Estimates for the GARCH Model on Jump-Filtered Returns for 20 Stocks and the S&P 500 Index (in bold) where the Coefficients are Represented by: $\tilde{R}_{t+1} - J_t = r + \lambda h_{t+1} + h_{t+1}^{1/2} z_{t+1}$, $h_{t+1} = \omega + \beta h_t + \alpha (z_t - \gamma h_t^{1/2})^2$ and $\mathcal{L}_{i,T}^*$ is the Log-Likelihood of the Sample for the Parameters at the Optimum

	AA	AXP	BA	CAT	CVX	DD	DIS	GE	HD	HPQ	IBM
λ	-1.861e+00 [2.465e-03]	-8.995e-02 [2.504e-03]	8.786e-02 [2.577e-03]	-4.376e-01 [2.617e-03]	-8.362e-02 [4.760e-03]	-1.014e+00 [2.814e-03]	-2.567e-01 [2.473e-03]	-8.532e-01 [2.659e-03]	-8.663e-01 [2.413e-03]	7.596e-01 [3.205e-03]	4.978e-01 [1.994e-03]
ω	5.490e-11 [8.899e-01]	2.334e-13 [5.571e+00]	1.647e-12 [5.064e+00]	7.510e-12 [1.960e+00]	4.749e-11 [1.155e+00]	6.865e-12 [2.778e+00]	2.878e-10 [1.122e+00]	3.316e-13 [3.825e+00]	6.483e-07 [1.839e-02]	1.419e-14 [1.325e+01]	1.206e-06 [1.568e-02]
β	9.678e-01 [2.354e-03]	9.191e-01 [1.922e-03]	9.310e-01 [1.688e-03]	9.500e-01 [2.514e-03]	9.170e-01 [2.258e-03]	9.294e-01 [1.656e-03]	9.107e-01 [1.788e-03]	9.221e-01 [1.968e-03]	9.134e-01 [2.139e-03]	9.727e-01 [3.672e-03]	8.291e-01 [1.468e-03]
α	1.310e-05 [2.780e-03]	1.711e-05 [2.205e-03]	1.856e-05 [1.870e-03]	4.619e-06 [2.033e-03]	9.416e-06 [2.290e-03]	1.644e-05 [1.952e-03]	2.912e-05 [1.886e-03]	4.874e-06 [1.788e-03]	9.248e-06 [2.207e-03]	6.326e-06 [4.541e-03]	2.922e-05 [1.274e-03]
γ	2.452e+01 [4.736e-04]	4.709e+01 [4.527e-04]	2.586e+01 [3.075e-04]	9.037e+01 [6.867e-04]	5.775e+01 [6.987e-04]	2.927e+01 [3.487e-04]	1.251e+01 [2.511e-04]	1.115e+02 [7.098e-04]	8.077e+01 [5.247e-04]	4.467e+01 [1.184e-03]	6.119e+01 [2.377e-04]
$\mathcal{L}_{i,T}^*$	8.580e+03	1.031e+04	1.128e+04	9.853e+03	9.133e+03	1.137e+04	1.153e+04	1.246e+04	1.113e+04	7.347e+03	1.179e+04
	JNJ	KO	MCD	PFE	PG	SPX	T	UTX	VZ	WMT	CJOW08
λ	-4.729e-01 [3.811e-03]	2.033e-02 [2.421e-03]	1.178e+00 [3.144e-03]	1.636e-03 [2.511e-03]	3.125e+00 [3.401e-03]	2.324e+00 [4.157e-03]	-1.488e+00 [2.635e-03]	-7.633e-01 [1.726e-03]	-1.099e+00 [4.117e-03]	-7.614e-02 [2.365e-03]	2.231e+00 [1.123e+00]
ω	8.409e-13 [3.812e+00]	3.256e-06 [8.237e-03]	1.726e-12 [2.600e+00]	3.296e-11 [1.493e+00]	2.201e-05 [1.660e-03]	8.847e-13 [1.321e+01]	1.567e-12 [3.880e+00]	4.194e-06 [1.114e-02]	5.551e-14 [7.473e+00]	3.163e-09 [2.770e-01]	2.101e-17 [1.120e-07]
β	9.132e-01 [1.589e-03]	6.784e-01 [1.095e-03]	9.667e-01 [2.124e-03]	9.138e-01 [1.803e-03]	5.029e-01 [1.020e-03]	8.212e-01 [1.353e-03]	9.479e-01 [2.267e-03]	5.503e-01 [1.061e-03]	9.677e-01 [2.851e-03]	7.961e-01 [9.985e-04]	9.012e-01 [4.678e-03]
α	9.247e-06 [1.702e-03]	3.883e-05 [1.029e-03]	6.031e-06 [1.912e-03]	2.702e-05 [1.936e-03]	8.656e-05 [8.065e-04]	4.306e-06 [1.266e-03]	1.420e-05 [2.468e-03]	3.815e-04 [1.190e-03]	6.189e-06 [2.759e-03]	7.109e-05 [1.039e-03]	3.317e-06 [1.380e-07]
γ	4.530e+01 [4.510e-04]	7.098e+01 [2.295e-04]	2.892e+01 [5.402e-04]	1.410e+01 [2.693e-04]	1.745e+01 [1.463e-04]	1.834e+02 [7.822e-04]	1.712e+01 [3.245e-04]	-1.683e+01 [9.109e-05]	1.844e+01 [6.350e-04]	-9.089e+00 [1.393e-04]	1.276e+02 [8.347e+00]
$\mathcal{L}_{i,T}^*$	1.388e+04	1.269e+04	1.253e+04	1.165e+04	1.260e+04	1.437e+04	1.223e+04	8.371e+03	9.803e+03	1.213e+04	-

Notes: In the last column in bold we report the coefficient estimates from Christoffersen et al. (2008), labeled CJOW08. Standard errors are computed from the information matrix formed by inverting the Hessian of the log-likelihood function at the optimum, $\mathcal{H}^* = [\partial^2 \mathcal{L}_{i,T}^* / \partial \theta_i \partial \theta_j]$ in the standard manner. The finite difference operator used to compute the Hessian is $\sqrt{\epsilon}$, where ϵ is the floating point relative accuracy. In contrast, the standard errors in CJOW08, probably due to computational issues, are computed using the outer product of the gradient $\nabla \mathcal{L}_{i,T}^*$ vector at the optimum. While robust to floating point errors in the finite differencing at the optimum, using the gradient does not asymptotically converge to the Cramer-Rao lower bound and hence the asymptotic parameter covariance matrix. A color version of this exhibit appears in the online edition of this article.

on a longer daily sample of unfiltered returns. The direct comparator is the column in bold, which shows our results for the jump-filtered daily returns on the S&P 500 index. It is also interesting to compare the point estimates from Heston and Nandi (2000), who computed $\lambda = 0.205$ compared to values that are about one order of magnitude higher for both Christoffersen et al. (2008) and our own estimates for the jump-filtered data. Similarly, the values of $\omega = 5.02e-6$, $\alpha = 1.32e-6$ and $\beta = 0.589$ appear very different from our own estimates and those suggested in Heston and Nandi (2000). Indeed, the long run volatilities suggested in Heston and Nandi (2000) appear radically different from those computed herein and in Christoffersen et al. (2008)—at between 8% and 10% for Heston and Nandi (2000) and between 13% and 17% for Christoffersen et al. (2008) and our jump-filtered results respectively for the S&P 500. Recall that the long-run variance is $\bar{h} = (\omega + \alpha) / (1 - \beta - \alpha\gamma^2)$ and using the “rule-of-16,” the annualized volatility is therefore $\sqrt{252\bar{h}}$. This discrepancy seems less likely due to the intercept ω , found to be $2.101e-17$ by Christoffersen et al. (2008) and $8.847e-13$ from the jump-filtered data, and might be more related to the estimation of β , which as both Christoffersen et al. (2008) and the jump-filtered returns indicate, is close to unity vis-à-vis the 0.589 suggested in Heston and Nandi (2000). Due to the requirements of the jump filtration step and data availability, our results run over the 1996–2014 period, a shorter sample than in Christoffersen et al. (2008).

Note that at the individual stock level, the variation in the parameters is quite high. IBM is an interesting example in that it appears to provide us with as close to an “ideal” stock as possible. For IBM the risk premium is $\lambda \approx -0.5$ and the long-run annualized volatility is 21%. For all stocks the standard errors indicate that the asymmetry γ parameters are high (in the high teens and sometimes in the 100s), but none are close to the 421 value reported in Heston and Nandi (2000). In fact Verizon (VZ) and Walmart (WMT) have small negative values for γ that are not statistically significant. The persistence parameter β is uniformly high, at above 0.9 for most stocks except Coca-Cola (KO) and Walmart but also uniformly below unity.

The results from Heston and Nandi (2000) for imputing the parameters of the model by nonlinear least squares from a small sample of options are even more at odds with the results found herein and as documented

in Christoffersen et al. (2008). While this is a relatively trivial exercise with our short maturity near-the-money option data, the results are extremely volatile and goodness of fit is several orders of magnitude away from our version of the Pan (2002) approach. This is available from the authors or can be computed very quickly using the Matlab code from this study. We therefore omit a detailed analysis from our discussion.

For the GARCH model (more than a decade since its introduction) we suggest that, while the GARCH option pricing framework may be closer to a structural model of options prices than the Black–Scholes model or other more primitive stochastic volatility models, the maximum likelihood parameter estimates still appear somewhat unstable, although far less unstable than the comparison between the results in Christoffersen et al. (2008) and Heston and Nandi (2000). Note that cross-sectionally the parameter estimates for individual stocks are very stable for the jump-filtered returns and indeed the jump-filtered returns do not exhibit substantial variation from the unfiltered returns (see the Online Supplement).

The CGARCH results presented in Exhibit 5 are again very consistent with our jump-filtered results and those found in Christoffersen et al. (2008). For the S&P 500 (center column in bold) we find λ to be highly consistent with the GARCH (similarly to Christoffersen et al. 2008). For our jump-filtered results we find that γ_2 , the asymmetry parameter in the component equation, is about 50% larger. We ascribe this variation to the 2007–2010 financial crisis period. The CGARCH model performs less consistently than the GARCH across the individual stocks, and the variation between the jump-filtered parameter estimates and the unfiltered estimates (Online Supplement), both individually and on average, is more pronounced. We will see from the option pricing performance that the CGARCH appears more acutely vulnerable to a lack of filtering jumps than the GARCH (possibly due to the component equation specification). Recall that we have restricted ourselves to the case when ρ is a free parameter, and in each case the point estimate of ρ is less than unity. Auxiliary likelihood ratio testing (not presented for space reasons) shows that for several cases, such as Boeing (BA), Caterpillar (CAT), and Dupont (DD), the restriction of $\rho = 1$ cannot be rejected at a 95% confidence level. Overall, the results indicate uniformly that the component persistence is

EXHIBIT 5

Maximum Likelihood Estimates for the CGARCH Model on Jump-Filtered Returns with Specification: $\tilde{R}_{t+1} - J_t = r + \lambda h_{t+1} + h_{t+1}^{1/2} z_{t+1}$ and $h_{t+1} = q_t + \beta(h_t - q_t) + \alpha((z_t - \gamma_1 h_t^{1/2})^2 - (1 + \gamma_1^2 q_t))$ where $q_t = \omega + \rho q_t + \varphi((z_t^2 - 1) - 2\gamma_2 h_t^{1/2} z_t)$

	AA	AXP	BA	CAT	CVX	DD	DIS	GE	HD	HPQ	IBM
λ	-1.859e+00 [1.007e-02]	-8.995e-02 [2.526e-02]	-1.443e-02 [8.075e-03]	-4.343e-01 [1.228e-01]	-4.492e-02 [1.470e-02]	-1.014e+00 [2.675e-02]	-3.658e-01 [7.919e-03]	-8.111e-01 [8.690e-03]	-8.708e-01 [1.993e-03]	7.727e-01 [1.075e-02]	4.978e-01 [6.209e-02]
β	9.524e-01 [3.879e-03]	9.160e-01 [8.242e-02]	9.050e-01 [5.189e-03]	9.530e-01 [6.064e-03]	8.966e-01 [6.161e-03]	9.016e-01 [3.027e-02]	8.758e-01 [1.588e-02]	8.393e-01 [5.602e-03]	9.012e-01 [2.202e-02]	6.646e-01 [1.003e-02]	9.652e-01 [3.333e-02]
α	1.136e-05 [1.909e-03]	1.656e-05 [3.432e-02]	1.447e-05 [3.264e-03]	3.237e-06 [8.512e-03]	1.928e-07 [3.334e-02]	2.538e-05 [2.719e-02]	1.488e-06 [2.426e-02]	7.498e-06 [4.478e-03]	6.593e-07 [1.004e-02]	3.527e-05 [3.859e-03]	2.657e-06 [1.688e-01]
γ_1	5.648e-07 [1.654e-06]	1.613e+02 [2.797e-02]	3.767e+01 [3.827e-06]	1.455e-11 [9.581e-06]	3.663e+03 [2.397e-03]	1.245e+02 [3.201e-02]	8.731e-10 [1.712e-05]	3.037e+01 [5.231e-06]	2.765e-10 [1.336e-05]	5.528e+00 [2.815e-06]	9.778e+02 [3.881e-02]
ω	4.098e-06 [3.081e-03]	2.089e-06 [1.740e-02]	1.806e-06 [6.187e-03]	2.826e-06 [5.789e-02]	3.602e-06 [3.909e-03]	4.762e-06 [7.697e-02]	2.443e-05 [2.574e-03]	2.002e-06 [2.504e-03]	6.416e-06 [2.113e-03]	2.053e-06 [4.654e-03]	8.681e-07 [1.384e-01]
ρ	9.916e-01 [3.990e-03]	9.900e-01 [2.741e-02]	9.946e-01 [6.101e-03]	9.907e-01 [3.336e-02]	9.810e-01 [3.780e-03]	9.900e-01 [2.652e-02]	9.297e-01 [2.677e-03]	9.933e-01 [2.650e-03]	9.894e-01 [2.160e-03]	9.949e-01 [4.455e-03]	9.900e-01 [1.777e-01]
ϕ	9.917e-06 [2.506e-03]	1.711e-05 [1.738e-02]	4.273e-06 [4.365e-03]	4.676e-06 [2.694e-02]	8.440e-06 [2.108e-03]	1.644e-05 [2.536e-02]	2.635e-05 [2.539e-03]	5.021e-06 [2.080e-03]	9.418e-06 [2.725e-03]	1.834e-06 [8.868e-03]	2.922e-05 [1.153e-01]
γ_2	6.810e+01 [2.883e-04]	5.690e-01 [2.298e-02]	1.077e+01 [3.202e-04]	8.604e+01 [3.671e-03]	2.191e-02 [6.301e-04]	4.614e+00 [4.685e-02]	1.309e+01 [1.407e-04]	6.734e+01 [3.463e-04]	8.459e+01 [2.933e-04]	1.226e+02 [2.211e-03]	2.955e-01 [1.888e-01]
$\mathcal{L}_{i,T}^*$	8.573e+03	2.990e+04	1.130e+04	9.860e+03	9.156e+03	1.872e+04	1.153e+04	1.253e+04	1.109e+04	7.384e+03	9.086e+03
	JNJ	KO	MCD	PFE	PG	SPX	T	UTX	VZ	WMT	CJOWO8
λ	-4.729e-01 [1.072e-02]	-6.977e-01 [9.059e-03]	1.200e+00 [1.145e-02]	1.636e-03 [4.333e-02]	3.125e+00 [2.881e-15]	2.321e+00 [2.729e-02]	-1.488e+00 [8.901e-03]	-7.633e-01 [1.208e-02]	-1.097e+00 [1.231e-02]	-7.614e-02 [1.251e-02]	2.092e+00 [7.729e-01]
β	9.019e-01 [1.361e-02]	8.393e-01 [1.870e-03]	9.538e-01 [1.299e-02]	9.684e-01 [3.450e-02]	3.609e-01 [3.426e-24]	7.426e-01 [3.781e-03]	9.513e-01 [5.970e-02]	4.395e-01 [9.094e-03]	9.522e-01 [1.206e-02]	9.593e-01 [3.975e-02]	6.437e-01 [2.759e-02]
α	3.327e-06 [1.020e-02]	3.167e-06 [3.366e-03]	1.260e-06 [8.243e-03]	1.020e-06 [4.791e-02]	1.225e-06 [3.304e-15]	3.712e-08 [3.449e-02]	4.132e-06 [1.303e-02]	1.301e-06 [1.831e-02]	3.667e-06 [9.530e-03]	2.727e-06 [1.452e-02]	1.580e-06 [2.430e-07]
γ_1	2.295e+02 [6.858e-03]	8.389e+02 [5.520e-05]	2.212e+02 [3.534e-05]	7.344e+02 [4.469e-02]	6.760e+02 [1.068e-24]	2.544e+04 [1.857e-02]	2.366e+02 [3.545e-02]	1.040e+03 [1.483e-02]	0.000e+00 [9.500e-06]	7.486e+02 [1.572e-02]	4.151e+02 [6.341e+01]
ω	1.673e-06 [1.682e-02]	2.003e-05 [1.359e-03]	1.068e-06 [2.769e-03]	6.541e-07 [2.660e-02]	4.073e-05 [7.145e-15]	2.042e-06 [2.134e-03]	1.046e-06 [8.608e-03]	1.367e-04 [1.415e-02]	1.988e-06 [5.202e-03]	1.216e-06 [6.003e-03]	8.208e-07 [7.620e-08]
ρ	9.900e-01 [1.215e-02]	9.384e-01 [1.321e-03]	9.950e-01 [2.235e-03]	9.900e-01 [8.787e-03]	9.900e-01 [1.338e-24]	9.843e-01 [2.276e-03]	9.900e-01 [1.572e-02]	9.900e-01 [9.317e-03]	9.913e-01 [4.789e-03]	9.900e-01 [4.566e-03]	9.896e-01 [9.630e-04]
ϕ	9.247e-06 [1.812e-02]	2.038e-05 [1.228e-03]	3.823e-06 [1.482e-03]	2.702e-05 [5.390e-02]	8.655e-05 [4.039e-24]	3.724e-06 [1.437e-03]	1.420e-05 [2.121e-02]	3.815e-04 [8.663e-03]	4.561e-06 [1.679e-03]	7.109e-05 [9.893e-03]	2.480e-06 [1.160e-07]
γ_2	8.784e+00 [1.415e-02]	3.486e-10 [3.578e-04]	1.000e-11 [5.029e-04]	7.203e-01 [6.039e-02]	2.745e+01 [8.941e-16]	1.056e+02 [4.856e-04]	4.224e+00 [2.152e-02]	9.130e-01 [9.948e-03]	6.895e+01 [8.839e-04]	1.431e+01 [2.780e-02]	6.324e+01 [5.300e+00]
$\mathcal{L}_{i,T}^*$	2.138e+04	1.271e+04	1.260e+04	1.193e+04	3.060e+03	1.442e+04	1.072e+03	5.871e+03	9.815e+03	7.139e+03	-

Notes: Standard errors are computed in the same manner as for the GARCH model. A color version of this exhibit appears in the online edition of this article.

high for both filtered and unfiltered returns, in line with Christoffersen et al. (2008).

It is worth stressing that our AR-Jump model with GARCH differs slightly from Christoffersen et al. (2008) in that we explicitly allow the jump intensity to vary stochastically. In Exhibit 6 we report the parameter estimates from our multistep estimation framework for the AR-Jump model with GARCH volatilities. Once again we report the estimates from Christoffersen et al. (2008) in bold in the last column. However, in this case they are not directly analogous for two reasons. First, our primary specification does not allow for jumps in variance. Indeed, this is a second maximum likelihood only estimation that is reported in the Online Supplement and corresponds to model #6 in our option comparison table. For brevity, we do not fully elucidate the results here. However, they are less consistent than the GARCH and CGARCH models with the equivalent results in Christoffersen et al. (2008). Second and more importantly, we allow the intensities to be autoregressive. Hence, two more parameters have to be identified rather than simply having a single long-run intensity.

We will present some interesting features of the results in Exhibit 6. However, it is important to note the speed advantage offered by estimating the model using the multistep approach outlined in the previous section. In contrast to the unfiltered MLE approach used to estimate the model parameters from daily unfiltered data, the MLE procedure, while iterative, only has to optimize the log-likelihood function for the jump risk premium. The preceding steps from the tick data are determined using descriptive statistics in a single step. The GARCH model used to extract the conditional variances is a five-parameter optimization and is also extremely fast. In contrast, full maximum likelihood estimation of the jump model requires an 11-dimensional optimization, and the computation of the optimal parameters is very slow. It is worth stressing that the log-likelihoods for the reported parameters (MS), while lightly lower than those reported for a fully optimized system (MLE), almost uniformly are lower than one-half of a chi-squared at 95%, i.e., $\mathcal{L}_{i,T}^{*,MLE} - \mathcal{L}_{i,T}^{*,MS} < 1/2\chi^2(1, 0.95)$. This suggests that the real data behaves in a statistically similar manner to the Monte Carlo study performed in the Online Supplement. Furthermore, the log-likelihoods, reported in Exhibit 6, are uniformly higher than the CGARCH and nested GARCH models, and in most cases substantially greater than $1/2\chi^2(6, 0.95)$ from the GARCH,

indicating that the joint restriction on the jump parameters is rejected. This is far more clear-cut than the likelihood ratio testing between the N and CGARCH cases (a restriction of three degrees of freedom). We also performed likelihood ratio tests to compare the AR versus constant jump intensity and found that the restriction was rejected each time, as is customary in the option pricing literature.

In summary, our results indicate that the maximum likelihood estimation of the models for filtered and unfiltered returns does change the fit to the underlying data, although not dramatically. We conclude this section by analyzing our out-of sample options results, for which we find substantial variation. Nevertheless, we want to reinforce the main theme of the present paper as the results are most likely driven by the inherent properties of the models and not due to substantial variation in the quality of fitting to the underlying data.

Out-of-Sample Options Results

Exhibits 7 and 8 present the root mean mean and median absolute errors, respectively. Recall that the loss functions here are relative valued, as they are reported from the squared differences in the natural logarithm of actual and predicted options prices. The bins correspond to those found in Exhibit 3. However, for brevity we present only three blocks of results on the central tendency of the out-of-sample pricing errors: First, those for the individual stocks as an aggregate; second, the results for the S&P 500 index, for comparison with Christoffersen et al. (2008) and Heston and Nandi (2000); and finally the results for IBM, which we chose as being a fairly archetypal stock from the in-sample model parameter estimates, and which has a sufficient number of options to provide a meaningful comparison to the aggregate stock data and the index.

The models are ordered using the summary outlined above. Therefore, #1 presents the option pricing errors for the GARCH model estimated on jump-filtered returns; #2 reports the option pricing errors for the CGARCH model estimated on jump-filtered returns; #3 presents the results for the AR-Jump model with GARCH variances estimated using the multistep model; #4 and #5 display the results for the GARCH and CGARCH model estimated on unfiltered data for comparison; and #6 presents the results for the out-of-sample option pricing errors for the ARJ-GARCH with

jumps in mean and variance, following Christoffersen et al. (2008) estimates by maximum likelihood. Recall that, for brevity, we only have reported the in-sample parameter estimates for the first three models. The remaining three model parameter estimates are available in the Online Supplement.

A useful approach to loss function analysis is to treat alternative methods separately and then compare the properties in light of the differences that each loss function is designed to detect. We start by reviewing the results from Exhibit 7, which reports the RMSE for the log values of the dollar option price. To convert these into dollar equivalents we compute $\exp(R-SE) - 1$.

Moving from the upper left in the exhibit, we can see that the six models perform very differently for relatively short maturity options. Exceptionally, the jump-filtered GARCH is consistently the best performer (recall that it was statistically not the best fit in terms of likelihood). None of the models appear to fully capture the Black–Scholes option-implied “smirk” pricing near-the-money options more effectively than out-of-the-money or heavily in-the-money options. The jump-filtered GARCH consistently fits the options better than the other models (including both of the GARCH with jumps models, #3 and #6). Following our discussion on the numerical stability of the CGARCH, we see that the jump-filtered and non-jump-filtered CGARCH offer consistently poor RMSE performance out-of-sample. In addition, moving down the table from 0–75-day maturity options, 75–100-day options and longer maturities, we see that the GARCH on the jump-filtered returns provides the lowest out-of-sample RMSE. Moving across the aggregate errors from all the stocks to the S&P 500 index options, we also can observe that the index options are uniformly priced more accurately than the options on the individual stocks. Notice that for the index the differentials in the pricing errors between the six models are now extremely modest. For instance, shorter maturity options from 0 to 75 days to maturity are almost uniformly smaller than a dollar for our candidate models (#1 to #3), with the minimum error being for the $S/K = 0.75$, for the GARCH of $\exp(0.284) - 1 \equiv 32.84$ cents and the maximum being jointly the GARCH and the CGARCH for heavily in-the-money $S/K > 1.5$ options, being $\exp(0.497) - 1 \equiv 64.38$ cents.

Moving down the center of Exhibit 7 for the S&P 500 index options, we can see that the pricing discrepancy (for all models) is mostly attributable to heavily

out-of-the-money long-maturity options $T - t > 250$ and $S/K < 0.75$. Noticeably, the CGARCH estimated on the unfiltered data appears to perform very poorly. The CGARCH on jump-filtered data, however, does not perform so poorly in the overall aggregate (center block, last row) and each of the models that either filter jumps or directly incorporate them into the specification perform well, albeit the overall best performer (by less than a cent) is the GARCH on unfiltered data.

The mean results for IBM paint a totally different picture. Recall that there are 22,641 contracts for IBM. Therefore, it is unlikely that the RMSEs are due to small-sample properties. The first point to note is that the AR-Jump model performs far better than the GARCH or the CGARCH estimated on jump-filtered returns. Indeed, for IBM the degree of fit for the multistage estimated AR-Jump GARCH model is exceptional. For instance, 75–100-day options (which contain one of the lowest sample counts at 49 for near-the-money options) exhibit a pricing accuracy of near $\exp(0.024) - 1 = 2.43$ cents for in-the-money contracts $1.025 < S/K \leq 1.25$. This performance is largely superior to the jumps in mean and variance model #6 and indeed for all the other models herein. For IBM, the AR-Jump model estimated stepwise from mixed-frequency data is very dominant, with a pricing error of less than 46 cents overall ($\exp(0.378) - 1 = 0.4594$).

RMSEs are very standard in the out-of-sample forecasting literature. Nonetheless, one issue frequently encountered is that they tend to be biased upward by a small number of dramatically poor-fitting forecast-realization pairs. Median absolute errors (or more commonly MAEs) provide a useful robust alternative. In Exhibit 8 we present the median absolute deviation (MAD). The first point to note is that for the overall results (reported in the last row), the MAE is consistently lower than the RMSE, indicating that large errors are indeed shifting the average higher. The stability issue for the CGARCH model is also remarkable, as this model considerably benefits from using the MAD rather than the RMSE for S&P 500 options.

It is interesting to note, however, that for the 203,587 options contracts on individual stocks the MAE does not deviate substantially from the RMSE for any of the models. This suggests that the pricing for a collection of individual stocks is a rather more complicated exercise than for an individual stock such as IBM or the S&P 500. Thus, in line with anecdotal evidence,

EXHIBIT 6

Multistep Maximum Likelihood Estimates for the AR-Jumps Model with GARCH Variances Using High-Frequency Data

	AA	AXP	BA	CAT	CVX	DD	DIS	GE	HD	HPQ	IBM
λ_z	-1.627e+00 [9.930e-02]	1.095e+00 [5.617e-02]	3.333e-01 [2.869e-02]	2.195e-01 [6.051e-02]	-6.052e-01 [-]	-7.258e-01 [3.405e-02]	-6.175e-01 [3.939e-02]	-5.640e-02 [4.569e-02]	-5.865e-01 [9.933e-02]	1.202e+00 [4.177e-02]	1.483e+00 [5.762e-02]
λ_y	-3.088e-03 [1.140e+02]	-4.535e-03 [6.299e-01]	-1.564e-03 [1.932e+01]	-3.989e-03 [1.812e+00]	-9.920e-31 [-]	-3.614e-03 [5.181e+00]	-1.937e-03 [4.355e+00]	-5.657e-03 [1.414e+00]	-8.025e-05 [6.812e+00]	-1.446e-03 [1.616e+01]	-2.648e-02 [1.369e+00]
ω	-3.384e-06 [4.668e-02]	-2.935e-06 [4.173e-02]	-3.726e-06 [7.292e-02]	-1.514e-06 [9.663e-02]	-2.152e-166 [-]	-3.600e-06 [2.455e-02]	-4.454e-07 [2.493e+00]	-2.195e-06 [3.888e-02]	-9.620e-09 [4.657e+00]	-1.966e-06 [6.171e-02]	-1.285e-08 [1.259e+02]
β	9.934e-01 [5.407e-02]	9.842e-01 [3.922e-02]	9.782e-01 [7.414e-02]	9.907e-01 [6.026e-02]	9.681e-01 [-]	9.882e-01 [7.765e-02]	9.326e-01 [2.411e-02]	9.831e-01 [3.315e-02]	9.954e-01 [8.460e-02]	9.909e-01 [5.975e-02]	9.089e-01 [1.029e-02]
α	5.083e-06 [2.569e-02]	4.300e-06 [3.307e-02]	8.275e-06 [4.642e-02]	2.071e-06 [4.637e-02]	3.197e-06 [-]	5.319e-06 [3.249e-02]	1.587e-05 [3.248e-02]	3.493e-06 [2.033e-02]	1.652e-08 [3.080e-02]	3.228e-06 [2.131e-02]	1.524e-05 [1.442e-02]
γ	-1.152e+01 [2.421e-06]	3.343e+01 [2.654e-06]	1.032e+01 [1.651e-06]	4.260e+01 [3.763e-06]	1.263e+01 [-]	1.814e+01 [1.781e-06]	1.331e+01 [1.130e-06]	4.719e+01 [2.331e-06]	4.718e+02 [4.850e-04]	2.244e+01 [3.819e-06]	3.133e+01 [1.006e-06]
χ_0	9.676e-04 [1.293e-01]	1.534e-03 [9.635e-02]	5.173e-03 [2.247e-01]	3.042e-03 [1.078e-01]	2.725e-03 [-]	8.085e-04 [1.803e-01]	1.219e-03 [4.593e-01]	2.409e-03 [8.169e-02]	1.272e-02 [8.399e-02]	1.390e-02 [4.492e-02]	7.427e-03 [1.696e+01]
ρ	9.832e-01 [2.460e+01]	9.887e-01 [6.515e-01]	7.747e-01 [6.795e-01]	9.529e-01 [1.821e-01]	9.865e-01 [-]	9.542e-01 [2.519e-01]	9.357e-01 [6.209e-01]	9.525e-01 [6.007e-01]	9.819e-01 [8.775e-02]	3.992e-01 [7.963e-01]	1.740e-01 [5.521e+02]
ϕ	2.510e-01 [1.406e+00]	2.209e-01 [4.953e-02]	3.625e-01 [1.164e-01]	1.911e-01 [9.819e-02]	2.958e-01 [-]	1.308e-01 [8.540e-02]	1.533e-01 [4.977e-01]	3.918e-01 [1.529e-01]	4.463e-01 [9.368e-02]	2.262e-01 [2.754e-02]	5.938e-05 [5.002e+03]
μ_j	-5.677e-02 [1.407e-02]	-3.997e-02 [9.380e-03]	-4.534e-02 [1.084e-02]	-3.889e-02 [4.394e-03]	-2.103e-02 [-]	-5.232e-02 [4.836e-03]	-6.642e-02 [2.011e-03]	-4.446e-02 [5.313e-03]	-1.764e-03 [1.067e-01]	-5.908e-02 [3.751e-03]	-6.625e-02 [4.100e-04]
σ_j	8.287e-04 [4.345e+01]	8.191e-04 [8.035e+01]	8.181e-04 [1.600e+02]	8.703e-04 [1.167e+01]	1.228e-06 [-]	9.018e-04 [1.865e+01]	8.841e-04 [7.317e+02]	8.586e-04 [3.508e+01]	1.827e-02 [4.634e-03]	8.450e-04 [1.274e+01]	1.369e-03 [2.666e+00]
$\mathcal{L}_{i,T}^*$	8.718e+03	1.050e+04	1.135e+04	9.953e+03	9.197e+03	1.147e+04	1.171e+04	1.260e+04	1.146e+04	7.483e+03	1.255e+04

(continued)

EXHIBIT 6 (continued)

Multistep Maximum Likelihood Estimates for the AR-Jumps Model with GARCH Variances Using High-Frequency Data

	JNJ	KO	MCD	PFE	PG	SPX	T	UTX	VZ	WMT	CJOW08
λ_z	1.000e+01 [3.160e+07]	7.492e-01 [1.303e-02]	1.152e+00 [3.981e-02]	3.233e-01 [3.461e-02]	2.793e+00 [2.722e-05]	3.120e+00 [1.925e-01]	-1.468e-02 [9.146e-02]	-1.193e+00 [4.195e-02]	-9.779e-01 [4.235e-02]	-2.420e-02 [1.339e-02]	2.793e+00 [1864e+00]
λ_y	-5.627e-04 [1.984e-01]	-1.107e-02 [2.130e+00]	-3.408e-03 [1.665e+03]	-3.115e-02 [3.411e-01]	-3.419e-03 [1.789e-03]	-1.746e-13 [1.548e+07]	-7.771e-03 [5.077e+00]	-1.210e-02 [1.949e+00]	-5.469e-03 [2.192e+03]	-5.209e-03 [1.959e+03]	-3.419e-03 [1.153e-02]
ω	-3.174e-07 [3.149e-01]	-3.247e-07 [1.217e+00]	-3.395e-06 [8.766e-03]	-2.618e-06 [3.003e-02]	-1.186e-06 [1.033e-02]	-8.269e-18 [1.940e+07]	-5.400e-06 [9.523e-03]	-3.385e-10 [6.258e+02]	-3.126e-06 [1.445e-02]	-7.225e-07 [9.999e-03]	-1.186e-06 [9.457e-08]
β	9.925e-01 [1.860e-01]	9.162e-01 [9.219e-03]	9.883e-01 [4.538e-02]	9.773e-01 [2.171e-02]	9.551e-01 [3.362e-07]	8.667e-01 [1.086e-01]	9.922e-01 [5.445e-02]	8.513e-01 [9.642e-03]	9.891e-01 [3.857e-02]	9.976e-01 [2.898e-02]	9.551e-01 [3.979e-03]
α	2.487e-07 [1.772e-01]	1.012e-05 [1.066e-02]	5.156e-06 [9.507e-03]	4.581e-06 [1.467e-02]	2.191e-06 [3.044e-10]	2.201e-06 [5.422e-02]	7.233e-06 [1.034e-02]	1.938e-05 [1.506e-02]	5.277e-06 [1.181e-02]	7.268e-07 [8.963e-03]	2.191e-06 [1.306e-07]
γ	-1.098e+02 [1.768e-04]	3.100e+01 [1.704e-06]	2.669e+01 [1.436e-06]	5.275e+01 [1.566e-06]	1.050e+02 [1.427e-10]	1.740e+02 [4.329e-05]	1.658e+01 [9.464e-07]	2.838e+01 [1.669e-06]	1.092e+01 [2.293e-06]	4.829e+01 [8.885e-07]	1.088e+02 [8.983e+00]
χ_0	4.766e-02 [7.440e-02]	5.802e-03 [1.917e-01]	1.146e-03 [1.655e+03]	4.202e-03 [1.226e-01]	1.000e-01 [8.833e-11]	8.798e-03 [2.397e-01]	1.260e-02 [6.217e+00]	6.037e-04 [3.430e-01]	1.483e-03 [2.291e+03]	1.956e-03 [2.138e+03]	9.845e-03* [2.100e-03]
ρ	9.522e-01 [1.087e-01]	5.413e-01 [6.789e-01]	4.856e-01 [1.235e+02]	1.424e-08 [7.620e+04]	9.000e-01 [1.308e-04]	9.828e-01 [1.017e-01]	7.750e-01 [1.746e+00]	9.769e-01 [3.472e-01]	4.885e-01 [1.938e+02]	6.188e-01 [9.881e+01]	
ϕ	7.811e-01 [6.686e-01]	1.480e-01 [4.670e-01]	7.379e-01 [3.296e+02]	6.854e-01 [6.614e-01]	5.000e-01 [3.532e-14]	3.428e-01 [2.154e-01]	6.090e-01 [8.571e+00]	3.668e-02 [2.876e-01]	5.988e-01 [1.362e+02]	7.527e-01 [8.952e+02]	
μ_j	-3.926e-04 [7.717e-01]	-4.687e-02 [4.690e-04]	-2.649e-04 [2.922e+03]	-9.962e-02 [1.622e-03]	-1.174e-02 [6.539e-10]	-6.430e-03 [6.640e-02]	-1.855e-03 [1.372e+00]	-3.660e-02 [1.515e-01]	-1.445e-04 [1.420e+03]	-1.739e-04 [1.809e+02]	-1.174e-02 [4.879e-03]
σ_j	9.401e-03 [7.479e-03]	1.583e-03 [2.890e+00]	7.833e-04 [2.654e+02]	1.608e-03 [8.364e+00]	7.834e-04 [8.620e-01]	7.629e-03 [1.154e-02]	7.978e-04 [4.174e+00]	1.015e-01 [2.913e-02]	7.868e-04 [1.068e+03]	7.885e-04 [2.174e+02]	2.799e-02 [2.048e-03]
$\mathcal{L}_{i,T}^*$	1.402e+04	1.338e+04	1.259e+04	1.192e+04	1.195e+04	1.450e+04	1.231e+04	9.711e+03	9.835e+03	1.218e+04	

Notes: The specification is: $\tilde{R}_{t+1} = r + \tilde{\lambda}_z h_t + \tilde{\lambda}_y \tilde{\chi}_t + h_t^{1/2} z_{t+1} + \gamma_t + \mu^* \tilde{\chi}_t$, where $\tilde{\chi}_t = \tilde{\chi}_0 + \rho \tilde{\chi}_{t-1} + \xi \eta_{t-1}$ and $h_{t+1} = \omega + \beta h_t + \alpha (z_t - \gamma h_t^{1/2})^2$ is the GARCH specification.

The Hessian of the maximized log-likelihood $\mathcal{L}_{0,T}^*$ is used to compute standard errors. Note that in the case of CVX the Hessian indicates that the standard errors vanish toward zero.

Bootstrap resampling and reconditioning of the Hessian suggest that the standard errors are in line with the other 20 data sets. Note that μ^* computed numerically from the pair of equations:

$$\mu^* = \mu_j - \nu \sigma_j^2 \text{ and } \bar{\chi}^* = \bar{\chi} e^{\frac{1}{2} \nu^2 \sigma_j^2 + \nu \mu_j}, \quad \bar{\chi} = \chi_0 / (1 - \rho). \text{ A color version of this exhibit appears in the online edition of this article.}$$

EXHIBIT 7

Mean Out-Of-Sample Option Pricing Error

	All Stocks						S&P 500 Index						IBM					
	#1	#2	#3	#4	#5	#6	#1	#2	#3	#4	#5	#6	#1	#2	#3	#4	#5	#6
0 < T - t < 75																		
0 < S/K ≤ 0.75	1.292	3.038	1.390	1.352	2.129	1.992	0.284	0.283	0.338	0.277	1.322	0.883	1.082	2.188	0.494	1.050	2.021	4.172
0.75 < S/K ≤ 0.975	1.000	2.479	1.644	1.082	1.667	1.605	0.408	0.407	0.391	0.400	1.586	0.905	0.746	1.688	0.299	0.708	1.800	3.890
0.975 < S/K ≤ 1.025	0.778	2.637	1.302	0.831	1.493	1.769	0.433	0.432	0.410	0.425	1.587	0.888	0.584	1.445	0.209	0.545	1.328	3.691
1.025 < S/K ≤ 1.25	0.806	2.924	1.527	0.848	1.509	1.883	0.491	0.491	0.465	0.485	1.673	1.029	0.514	1.318	0.175	0.475	1.362	3.575
1.25 < S/K ≤ 1.50	0.799	2.609	1.982	0.811	1.483	1.808	0.497	0.497	0.473	0.491	1.674	1.057	0.409	1.149	0.122	0.371	1.753	3.404
S/K > 1.5	1.027	2.727	1.405	1.133	1.944	1.700	0.324	0.323	0.326	0.316	1.272	0.723	0.602	1.340	0.263	0.576	1.690	3.408
75 < T - t ≤ 100																		
0 < S/K ≤ 0.75	0.967	2.246	1.477	1.061	1.961	2.104	0.224	0.218	0.285	0.207	1.005	0.706	0.288	0.959	0.068	0.252	2.458	3.248
0.75 < S/K ≤ 0.975	0.922	2.915	2.019	1.035	1.709	1.803	0.244	0.240	0.281	0.230	1.194	0.677	0.236	0.848	0.045	0.201	1.340	3.073
0.975 < S/K ≤ 1.025	0.980	2.831	1.706	1.142	1.850	1.812	0.266	0.263	0.294	0.254	1.208	0.758	0.237	0.858	0.047	0.203	3.002	3.095
1.025 < S/K ≤ 1.25	1.135	3.175	1.616	1.246	2.117	1.855	0.214	0.210	0.253	0.200	1.147	0.599	0.207	0.758	0.036	0.174	1.190	2.879
1.25 < S/K ≤ 1.50	1.145	2.777	1.642	1.232	2.051	1.794	0.259	0.255	0.307	0.246	1.150	0.760	0.167	0.682	0.024	0.137	1.321	2.791
S/K > 1.5	1.091	2.658	1.412	1.176	1.997	1.838	0.205	0.201	0.258	0.192	1.058	0.636	0.188	0.725	0.041	0.160	1.274	2.826
100 < T - t ≤ 150																		
0 < S/K ≤ 0.75	1.242	2.573	1.709	1.345	1.649	1.781	0.128	0.122	0.271	0.111	1.108	0.303	1.363	2.557	0.671	1.340	2.038	4.351
0.75 < S/K ≤ 0.975	0.910	2.506	1.472	0.994	1.545	1.658	0.226	0.221	0.260	0.211	1.266	0.596	0.886	1.864	0.392	0.854	1.776	3.942
0.975 < S/K ≤ 1.025	1.006	2.665	2.071	1.155	1.743	1.766	0.250	0.244	0.279	0.233	1.323	0.623	0.737	1.650	0.306	0.702	1.319	3.842
1.025 < S/K ≤ 1.25	1.063	3.021	1.787	1.152	1.888	1.764	0.279	0.275	0.322	0.265	1.216	0.789	0.677	1.545	0.274	0.643	1.801	3.698
1.25 < S/K ≤ 1.50	1.085	2.738	1.724	1.152	1.872	1.838	0.246	0.243	0.278	0.235	1.079	0.711	0.598	1.400	0.232	0.564	1.321	3.506
S/K > 1.5	1.055	2.701	1.770	1.162	1.823	1.813	0.196	0.192	0.245	0.184	1.060	0.613	0.560	1.358	0.207	0.525	1.725	3.510
150 < T - t ≤ 250																		
0 < S/K ≤ 0.75	1.502	2.777	1.871	1.533	2.084	1.632	0.553	0.552	0.490	0.545	1.858	0.970	2.259	4.128	1.188	2.345	3.836	2.712
0.75 < S/K ≤ 0.975	1.145	2.884	1.384	1.257	1.992	1.301	0.519	0.520	0.485	0.514	1.771	1.059	1.349	2.647	0.728	1.357	2.657	2.482
0.975 < S/K ≤ 1.025	0.892	2.873	1.237	0.968	1.683	1.395	0.474	0.474	0.452	0.468	1.599	1.049	1.007	2.113	0.530	0.992	2.897	2.484
1.025 < S/K ≤ 1.25	0.943	2.406	1.734	0.927	1.693	1.668	0.360	0.358	0.367	0.350	1.444	0.853	1.257	2.493	0.627	1.237	2.153	3.389
1.25 < S/K ≤ 1.50	0.886	2.633	1.571	0.920	1.570	1.671	0.335	0.333	0.348	0.326	1.335	0.842	0.972	2.100	0.443	0.940	1.927	3.406
S/K > 1.5	1.035	2.864	1.458	1.133	1.818	1.716	0.397	0.396	0.385	0.389	1.477	0.875	0.669	1.643	0.271	0.632	1.941	3.111
T - t > 250																		
0 < S/K ≤ 0.75	1.988	4.528	1.757	2.312	3.373	1.603	2.919	3.001	2.986	3.079	5.205	3.861	2.927	5.369	1.456	3.079	4.982	2.835
0.75 < S/K ≤ 0.975	1.413	2.994	1.526	1.498	2.433	1.470	1.120	1.134	0.988	1.142	2.494	1.495	1.652	3.216	0.870	1.659	2.843	3.122
0.975 < S/K ≤ 1.025	0.913	2.174	1.279	0.941	1.737	1.442	1.008	1.028	0.820	1.037	2.484	1.248	1.077	2.338	0.544	1.066	2.215	3.174
1.025 < S/K ≤ 1.25	0.535	1.941	1.500	0.585	1.288	1.403	1.007	1.025	0.818	1.035	2.462	1.153	0.568	1.568	0.218	0.529	2.036	2.929
1.25 < S/K ≤ 1.50	0.655	1.962	1.791	0.655	1.164	1.421	0.898	0.912	0.829	0.920	2.156	1.175	0.655	2.002	0.478	0.924	2.310	2.492
S/K > 1.5	0.744	1.789	1.492	0.681	1.236	1.221	1.004	1.029	0.943	1.049	2.194	1.305	0.168	0.710	0.079	0.160	1.490	1.813
Mean	1.032	2.702	1.609	1.111	1.817	1.685	0.526	0.530	0.522	0.528	1.647	0.972	0.822	1.823	0.378	0.814	2.061	3.229

EXHIBIT 8

Median Out-Of-Sample Option Pricing Error

	All Stocks						S&P 500 Index						IBM					
	Mod #1	Mod #2	Mod #3	Mod #4	Mod #5	Mod #6	Mod #1	Mod #2	Mod #3	Mod #4	Mod #5	Mod #6	Mod #1	Mod #2	Mod #3	Mod #4	Mod #5	Mod #6
0 < T - t < 75																		
0 < S/K ≤ 0.75	1.282	1.822	1.230	1.337	2.123	1.991	0.099	0.092	0.136	0.082	1.114	0.221	1.142	2.321	0.495	1.103	1.774	4.223
0.75 < S/K ≤ 0.975	0.891	1.289	1.185	1.006	1.486	1.672	0.316	0.311	0.318	0.298	1.494	0.713	0.796	1.805	0.300	0.752	1.774	3.943
0.975 < S/K ≤ 1.025	0.557	0.988	1.122	0.576	1.056	1.705	0.297	0.290	0.309	0.277	1.441	0.658	0.647	1.548	0.220	0.603	1.608	3.842
1.025 < S/K ≤ 1.25	0.508	0.909	1.056	0.464	1.048	1.844	0.355	0.351	0.353	0.339	1.580	0.851	0.556	1.408	0.179	0.513	1.426	3.618
1.25 < S/K ≤ 1.50	0.749	0.913	1.118	0.714	1.188	1.806	0.362	0.358	0.360	0.344	1.600	0.889	0.440	1.217	0.124	0.398	1.275	3.436
S/K > 1.5	0.735	1.137	1.091	0.806	1.513	1.655	0.195	0.188	0.223	0.175	1.137	0.438	0.366	1.098	0.089	0.325	1.125	3.383
75 < T - t ≤ 100																		
0 < S/K ≤ 0.75	0.969	1.613	1.371	1.061	1.909	2.074	0.223	0.217	0.280	0.206	1.010	0.701	0.294	0.994	0.062	0.256	1.008	3.271
0.75 < S/K ≤ 0.975	0.756	1.326	1.335	0.871	1.439	1.784	0.236	0.231	0.277	0.219	1.156	0.657	0.242	0.829	0.039	0.205	0.898	3.004
0.975 < S/K ≤ 1.025	0.916	1.429	1.366	1.102	1.664	1.790	0.220	0.214	0.273	0.203	1.216	0.610	0.237	0.890	0.036	0.201	1.774	3.206
1.025 < S/K ≤ 1.25	1.120	1.844	1.344	1.225	2.013	1.814	0.183	0.176	0.225	0.162	1.179	0.453	0.187	0.721	0.021	0.155	0.043	2.816
1.25 < S/K ≤ 1.50	1.145	1.835	1.392	1.218	2.052	1.780	0.256	0.251	0.300	0.240	1.190	0.797	0.154	0.648	0.015	0.125	1.774	2.711
S/K > 1.5	1.101	1.850	1.257	1.193	1.961	1.771	0.192	0.185	0.245	0.174	1.068	0.529	0.133	0.613	0.008	0.104	1.774	2.695
100 < T - t ≤ 150																		
0 < S/K ≤ 0.75	1.174	1.520	1.562	1.282	1.567	1.389	0.146	0.138	0.165	0.125	1.208	0.308	1.339	2.722	0.593	1.305	1.774	4.065
0.75 < S/K ≤ 0.975	0.720	1.003	1.084	0.834	1.315	1.511	0.184	0.176	0.219	0.163	1.287	0.437	0.586	1.445	0.198	0.543	1.608	3.781
0.975 < S/K ≤ 1.025	0.976	1.662	1.079	1.125	1.848	1.603	0.228	0.222	0.273	0.209	1.312	0.568	0.646	1.616	0.206	0.603	0.946	3.700
1.025 < S/K ≤ 1.25	0.900	1.565	1.173	0.988	1.717	1.720	0.264	0.259	0.312	0.247	1.165	0.809	0.558	1.338	0.171	0.515	1.774	3.527
1.25 < S/K ≤ 1.50	1.078	1.668	1.511	1.112	1.931	1.859	0.219	0.213	0.262	0.202	1.118	0.632	0.525	1.343	0.144	0.483	1.774	3.516
S/K > 1.5	0.967	1.645	1.535	1.084	1.677	1.789	0.136	0.129	0.182	0.117	1.092	0.353	0.572	1.513	0.184	0.529	1.618	3.632
150 < T - t ≤ 250																		
0 < S/K ≤ 0.75	1.050	1.467	1.353	1.078	1.508	1.191	0.628	0.626	0.515	0.615	1.942	1.014	2.153	4.241	0.820	2.167	1.774	2.155
0.75 < S/K ≤ 0.975	0.914	1.182	1.117	1.048	1.667	1.151	0.393	0.388	0.360	0.375	1.672	0.850	0.535	1.413	0.170	0.491	1.774	2.021
0.975 < S/K ≤ 1.025	0.724	1.162	0.967	0.772	1.379	1.120	0.436	0.435	0.422	0.427	1.568	1.117	0.422	1.245	0.112	0.381	1.774	2.243
1.025 < S/K ≤ 1.25	0.610	1.144	1.065	0.582	1.166	1.426	0.257	0.251	0.296	0.239	1.325	0.698	1.091	2.563	0.508	1.052	1.774	2.981
1.25 < S/K ≤ 1.50	0.579	1.153	1.208	0.575	1.290	1.526	0.224	0.218	0.278	0.207	1.213	0.672	0.797	2.103	0.330	0.755	1.774	3.048
S/K > 1.5	0.955	1.628	1.263	1.059	1.638	1.706	0.289	0.285	0.311	0.273	1.407	0.744	0.497	1.620	0.164	0.452	1.404	2.640
T - t > 250																		
0 < S/K ≤ 0.75	1.923	3.259	1.489	2.276	3.312	1.364	2.469	2.514	2.372	2.548	4.499	3.662	3.060	5.595	1.498	3.214	4.881	2.262
0.75 < S/K ≤ 0.975	1.198	1.921	1.257	1.304	2.080	1.108	0.335	0.329	0.306	0.316	1.690	0.609	1.465	2.983	0.720	1.442	2.120	2.400
0.975 < S/K ≤ 1.025	0.751	1.410	0.875	0.767	1.487	1.160	0.472	0.470	0.419	0.457	1.795	0.905	0.835	2.045	0.343	0.792	1.774	2.561
1.025 < S/K ≤ 1.25	0.415	0.936	0.721	0.443	1.114	1.104	0.353	0.348	0.325	0.335	1.604	0.711	0.519	1.503	0.167	0.474	1.362	2.369
1.25 < S/K ≤ 1.50	0.439	0.806	1.611	0.437	0.823	1.095	0.195	0.188	0.220	0.175	1.197	0.441	0.245	0.993	0.045	0.206	0.878	1.926
S/K > 1.5	0.609	0.639	1.132	0.508	0.916	0.917	0.099	0.094	0.176	0.087	0.834	0.341	0.073	0.580	0.009	0.049	0.449	1.363
Mean	0.891	1.425	1.229	0.962	1.597	1.548	0.343	0.339	0.358	0.328	1.438	0.747	0.704	1.699	0.266	0.674	1.583	3.011

our results suggest that for certain individual stocks, options contracts are extraordinarily difficult to price. For the S&P 500 models #1 to #4 and #6 are all very consistent. However, the CGARCH estimated on the unfiltered returns performs relatively poorly (similarly to the RMSE case) for most contracts' moneyness and maturity and while most of the models perform badly for heavily out-of-the-money $S/K < 0.75$ and long-maturity $T - t > 250$ contracts, the CGARCH on unfiltered returns performs particularly poorly out-of-sample.

Remarkably, we note that for IBM #3, the multi-stage AR-Jump model estimated from mixed-frequency data provides the best fit by a considerable margin across most contract types. The MAE for out-of-sample contract pricing drops below one cent for heavily in-the-money ($S/K > 1.5$) and intermediate maturity $75 < T - t < 100$ options, suggesting that for IBM the AR-Jump is capturing out-of-sample the option smirk. This performance is replicated for long-maturity smirks, where the $S/K > 1.5$, $T - t > 250$ MAE is also less than one cent. Of course, as the maturity lengthens, the convexity of the implied volatility smile decreases. Nevertheless, long-term equity options such as LEAPS are generally very difficult to correctly price for heavily out-of-the-money and heavily in-the-money options.

CONCLUSIONS

This paper has introduced a new method for estimating the parameters for a jump GARCH model using mixed-frequency data to disentangle the high-speed jumps for continuous volatility, and combining these with jump premiums from near-the-money, short-maturity options. We then estimated this model in sample using tick data for the S&P 500 and a selection of 20 stocks from the DOW 30 components for the 1996–2010 period. We have collected a very large sample of options from 2010 onward, and estimated the out-of-sample fit for the predicted options prices for our jump model versus alternative GARCH models estimated using jump-filtered and unfiltered daily returns. Our results are mixed. For IBM stock options, our jump model outperforms all others by more than an order of magnitude for many maturity moneyness combinations. However, for the S&P 500, the GARCH model estimated on jump-filtered returns appears to dominate. Nevertheless, the median performance of each of the models is very close, except for the CGARCH

model that contains a substantial persistence equation. The CGARCH performs consistently poorly across all our test cases, and we partly attribute this issue to inherent instabilities in the pricing mechanism. However, we also find that not filtering jumps from the continuous component appears to be a major factor for the deterioration of the CGARCH pricing performance. Our results are derived from a very large test bed of tick data and options contracts. Remarkably, this is a substantial innovation over previous research that has typically focused on calibration and pricing performance relative to much smaller options data sets. Indeed, for options requiring Monte Carlo estimation of AR-Jump models this appears to be the largest out-of-sample exercise of its type.

The heterogeneity of the results, particularly in aggregate for individual stocks, illustrates the inherent difficulty in correctly pricing options through a structural model. However, the performance of the AR-Jump model for a large cross-section of IBM call options paves the way for structural models with jumps and GARCH or stochastic volatilities that match the performance of options models directly calibrated to the implied volatility surface. Finally, we have presented a substantial body of evidence to suggest that the CGARCH model does not provide numerically stable options prices, and the use of this model to forecast options prices out of sample may be ill advised.

REFERENCES

- Aït-Sahalia, Y. 2004. "Disentangling Diffusion from Jumps." *Journal of Financial Economics* 74 (3): 487–528.
- Aït-Sahalia, Y., and J. Jacod. 2009a. "Estimating the Degree of Activity of Jumps in High Frequency Financial Data." *The Annals of Statistics* 37 (5A): 2202–2244.
- . 2009b. "Testing for Jumps in a Discretely Observed Process." *The Annals of Statistics* 37 (1): 184–222.
- Albrecher, H., P. Mayer, W. Schoutens, and J. Tistaert. 2006. "The Little Heston Trap." *Wilmott* 1: 83–92.
- Andersen, T. G., L. Benzoni, and J. Lund. 2002. "Empirical Investigation of Continuous-Time Equity Return Models." *The Journal of Finance* 57 (3): 1239–1284.
- Andersen, T. G., T. Bollerslev, F. X. Diebold, and P. Labys. 2003. "Modeling and Forecasting Realized Volatility." *Econometrica* 71 (2): 579–625.

- Babaoglu, K., P. Christoffersen, S. Heston, and K. Jacobs. 2017. "Option Valuation with Volatility Components, Fat Tails, and Nonmonotonic Pricing Kernels." *The Review of Asset Pricing Studies* 8 (2): 183–231.
- Bajgrowicz, P., and O. Scaillet. 2009. "We Propose a Technique to Avoid Spurious Detections of Jumps in High-frequency Data via an Explicit Thresholding on Available Test Statistics." Swiss Finance Institute Research Paper Series 11–36, Swiss Finance Institute.
- Bakshi, G., C. Cao, and Z. Chen. 1997. "Empirical Performance of Alternative Option Pricing Models." *Journal of Finance* 52 (5): 2003–2049.
- Ballotta, L., G. Deelstra, and G. Rayée. 2017. "Quanto Implied Correlation in a Multi-lévy Framework." Forthcoming, *European Journal of Operational Research* 260 (3): 1181–1199.
- Barndorff-Nielsen, O., and N. Shephard. 2004. "Power and Bipower Variation with Stochastic Volatility and Jumps." *Journal of Financial Econometrics* 2 (1): 1–48.
- . 2006. "Econometrics of Testing for Jumps in Financial Economics Using Bipower Variation." *Journal of Financial Econometrics* 4 (1): 1–30.
- Bates, D. S. 1996. "Jumps and Stochastic Volatility: Exchange Rate Processes Implicit in Deutsche Mark Options." *Review of Financial Studies* 9 (1): 69–107.
- . 2000. "'87 Crash Fears in the S&P 500 Futures Option Market." *Journal of Econometrics* 94 (1–2): 181–238.
- . 2006. "Maximum Likelihood Estimation of Latent Affine Processes." *Review of Financial Studies* 19 (3): 909–965.
- Bollerslev, T. 1986. "Generalized Autoregressive Conditional Heteroskedasticity." *Journal of Econometrics* 31 (3): 307–327.
- Broadie, M., M. Chernov, and M. Johannes. 2007. "Model Specification and Risk Premiums: Evidence from Futures Options." *Journal of Finance* 62 (3): 1453–1490.
- Carr, P., and L. Wu. 2003. "The Finite Moment Log Stable Process and Option Pricing." *Journal of Finance* 58 (2): 753–777.
- Chan, W. H., and J. M. Maheu. 2002. "Conditional Jump Dynamics in the Stock Market Returns." *Journal of Business and Economic Statistics* 20 (3): 377–389.
- Chernov, M., A. R. Gallant, E. Ghysels, and G. Tauchen. 2003. "Alternative Models for Stock Price Dynamics." *Journal of Econometrics* 116 (1–2): 225–257.
- Christoffersen, P., S. Heston, and K. Jacobs. 2006. "Option Valuation with Conditional Skewness." *Journal of Econometrics* 131 (1–2): 253–284.
- Christoffersen, P., and K. Jacobs. 2004. "The Importance of the Loss Function in Option Valuation." *Journal of Financial Economics* 72 (2): 291–318.
- Christoffersen, P., K. Jacobs, B. Feunou, and N. Meddahi. 2012. "The Economic Value of Realized Volatility: Using High-Frequency Returns for Option Valuation." SSRN Working Paper.
- Christoffersen, P., K. Jacobs, C. Ornathanalai, and Y. Wang. 2008. "Option Valuation with Long-Run and Short-Run Volatility Components." *Journal of Financial Economics* 90 (3): 272–297.
- Corsi, F., N. Fusari, and D. La Vecchia. 2013. "Realizing Smiles: Options Pricing with Realized Volatility." *Journal of Financial Economics* 107 (2): 284–304.
- Durham, G., J. Geweke, and P. Ghosh. 2015. "A Comment on Christoffersen, Jacobs, and Ornathanalai (2012), Dynamic Jump Intensities and Risk Premiums: Evidence from S&P 500 Returns and Options." *Journal of Financial Economics* 115 (1): 210–214.
- Engle, R. F. 1982. "Autoregressive Conditional Heteroskedasticity with Estimates of the Variance of UK Inflation." *Econometrica* 50 (4): 987–1008.
- Engle, R. F., and G. G. J. Lee. 1999. "A Permanent and Transitory Component Model of Stock Return Volatility." Technical report, In: Engle, R., and H. White (Eds.), "Cointegration, Causality, and Forecasting: A Festschrift in Honor of Clive W. J. Granger." New York, NY: Oxford University Press.
- Eraker, B. 2004. "Do Stock Prices and Volatility Jump? Reconciling Evidence from Spot and Option Prices." *Journal of Finance* 59 (3): 1367–1403.
- Eraker, B., M. Johannes, and N. Polson. 2003. "The Impact of Jumps in Volatility and Returns." *Journal of Finance* 53 (3): 1269–1300.

- Fuertes, A., and J. Olmo. 2012. "The Role of High-Frequency Prices, Long Memory and Jumps for Value-At-Risk Prediction." Technical report, Cass Business School Working Paper, available at SSRN: <http://ssrn.com/abstract=2070018>.
- Giacomini, R., and B. Rossi. 2010. "Forecast Comparisons in Unstable Environments." *Journal of Applied Econometrics* 25 (4): 595–620.
- Harvey, A. C., and C. Fernandes. 1989. "Time Series Models for Count or Qualitative Observations." *Journal of Business & Economic Statistics* 7 (4): 407–417.
- Heinen, A. 2003. "Modelling Time Series Count Data: An Autoregressive Conditional Poisson Model." CORE discussion paper. Louvain-la-Neuve, France: CORE.
- 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.
- Heston, S. L., and S. Nandi. 2000. "A Closed-Form GARCH Option Valuation Model." *Review of Financial Studies* 13 (3): 585–625.
- Huang, X., and G. Tauchen. 2005. "The Relative Contribution of Jumps to Total Price Variance." *Journal of Financial Econometrics* 3 (4): 456–499.
- Jacod, J., and V. Todorov. 2009. "Testing for Common Arrival of Jumps in Discretely-Observed Multidimensional Processes." *Annals of Statistics* 37 (4): 1792–1838.
- Jorion, P. 1998. "On Jump Processes in the Foreign Exchange and Stock Markets." *Review of Financial Studies* 1 (4): 427–445.
- Maheu, J. M., and T. H. McCurdy. 2004. "News Arrival, Jump Dynamics, and Volatility Components for Individual Stock Returns." *Journal of Finance* 20 (2): 755–793.
- Merton, R. 1976. "Option Pricing when Underlying Stock Returns Are Discontinuous." *Journal of Financial Economics* 3: 125–144.
- Mykland, P., Y. Aït-Sahalia, and L. Zhang. 2005. "A Tale of Two Time Scales: Determining Integrated Volatility with Noisy High-Frequency Data." *Journal of the American Statistical Association* 100 (472): 1394–1411.
- Mykland, P. A., and L. Zhang. 2009. "The Econometrics of High Frequency Data." Technical report, University of Illinois at Chicago.
- Palm, F. C., and P. J. G. Vlaar. 1993. "The Message in Weekly Exchange Rates in the European Monetary System: Mean Reversion, Conditional Heteroscedasticity, and Jumps." *Journal of Business & Economic Statistics* 11 (3): 351–360.
- Pan, J. 2002. "The Jump-Risk Premia Implicit in Options: Evidence from an Integrated Time-Series Study." *Journal of Financial Economics* 63 (1): 3–50.
- Roberts, H. 1959. "Stock-Market 'Patterns' and Financial Analysis: Methodological Suggestions." *Journal of Finance* 14 (1): 1–10.
- Scott, L. O. 1997. "Pricing Stock Options in a Jump-Diffusion Model with Stochastic Volatility and Interest Rates: Applications of Fourier Inversion Methods." *Mathematical Finance* 7 (4): 413–424.
- Tauchen, G., and V. Todorov. 2010. "Activity Signature Functions for High-Frequency Data Analysis." *Journal of Econometrics* 154 (2): 125–138.
- Tauchen, G., and H. Zhou. 2011. "Realized Jumps on Financial Markets and Predicting Credit Spreads." *Journal of Econometrics* 160 (1): 102–118.

To order reprints of this article, please contact David Rowe at d.rowe@pageantmedia.com or 646-891-2157.

ADDITIONAL READING

An Improved Estimation Method for a Family of GARCH Models

PASCAL LÉTOURNEAU

The Journal of Derivatives

<https://jod.pm-research.com/content/27/1/67>

ABSTRACT: This article proposes an improved estimation and calibration method to a family of GARCH models. The suggested method fixes one parameter such that the unconditional kurtosis of the model matches the sample kurtosis. An empirical analysis using Engle and Ng's (1993) NGARCH(1,1) model shows that the method dominates previous estimation methods in multiple ways. The optimization problem is simplified and made less sensitive to initial values. The optimization time, both when estimating based on historical returns and calibrating to option prices, is reduced by roughly 50%. The in-sample fit is barely affected, while the option pricing, both in sample and out of sample, is improved.

VIX Futures Pricing with Affine Jump-GARCH Dynamics and Variance-Dependent Pricing Kernels

XINGLIN YANG, PENG WANG, AND JI CHEN

The Journal of Derivatives

<https://jod.pm-research.com/content/27/1/110>

ABSTRACT: *Volatility Index (VIX) futures are among the most actively traded contracts at the Chicago Board Options Exchange, in response to the growing need for protection against volatility risk. The authors develop a new class of discrete-time and closed-form VIX futures pricing models, in which the S&P 500 returns follow the time-varying infinite-activity Normal Inverse Gaussian (NIG) and finite-activity compound Poisson (CP) jump-GARCH models, and which are risk-neutralized by the variance-dependent pricing kernel used by Christoffersen et al. (2013). They estimate these models using several data sets, including the S&P 500 returns, VIX Index, and VIX futures. The empirical results indicate that the time-varying NIG and CP jump-GARCH models significantly outperform the Heston-Nandi (HN) GARCH model in asset returns fitting and VIX futures pricing.*

Ensuring More Is Better: On the Simultaneous Application of Stock and Options Data to Estimate the GARCH Options Pricing Model

CHARLES CHANG, HUNG-WEN CHENG,
AND CHENG-DER FUH

The Journal of Derivatives

<https://jod.pm-research.com/content/26/1/7>

ABSTRACT: *The most common approach in fitting option pricing models to market data is first to make an assumption about the underlying asset's returns process and then develop an option pricing model for that process that is tested against market option prices. The returns process is estimated from historical data, option values are computed, and then compared against a cross-section of prices from the options market. Unfortunately, this often does not work well, and plainly it is inefficient in its use of the data. However, efforts to*

combine returns data from the asset market and prices from the options market into a single estimation have also not had much success. In this article, Chang, Cheng, and Fuh propose a new procedure to combine data from both markets in the estimation, in which options are assumed to be subject to random pricing noise relative to model values. The additional slack gives the estimator better ability to match prices in both markets. The article contrasts the performance of the full model approach with an approach that only uses stock prices or options prices to fit an option pricing model based on an underlying GARCH process. The value of the combined approach is demonstrated both theoretically as an asymptotic result in the model and also in a Monte Carlo simulation.