AARHUS UNIVERSITY
BUSINESS AND SOCIAL SCIENCES
DEPARTMENT OF ECONOMICS AND BUSINESS

CREATES
Center for Research in Econometric Analysis of Time Series

CREATES Research Paper 2011-34

American Option Pricing with Discrete and Continuous Time Models: An Empirical Comparison

Lars Stentoft

Department of Economics and Business
Aarhus University
Bartholins Allé 10
DK-8000 Aarhus C
Denmark

Email: oekonomi@au.dk
Tel: +45 8942 1610

American Option Pricing with Discrete and Continuous Time Models: An Empirical Comparison*

Lars Stentoft†

HEC Montréal, CIRANO, CIRPEÉ, and CREATES

September 25, 2011

Abstract

This paper considers discrete time GARCH and continuous time SV models and uses these for American option pricing. We first of all show that with a particular choice of framework the parameters of the SV models can be estimated using simple maximum likelihood techniques. Hence the two types of models can be implemented in an internally consistent manner. We then perform a Monte Carlo study to examine their differences in terms of option pricing, and we study the convergence of the discrete time option prices to their implied continuous time values. The results show that there are differences between the two models, though the discrete time GARCH prices converge quickly to the continuous time SV values. Finally, a large scale empirical analysis using individual stock options and options on an index is performed comparing the estimated prices from discrete time models to the corresponding continuous time model prices. The results show that, while the overall differences in performance are small, for the in the money put options on individual stocks the continuous time SV models do generally perform better than the discrete time GARCH specifications.

JEL Classification: C22, C53, G13

Keywords: American Options, Augmented GARCH, Least Squares Monte Carlo, Stochastic Volatility.

*The author thanks participants and the discussant at the 2008 Northern Finance Association Meeting as well as seminar participants at CREATES and HEC Montréal for valuable comments and suggestions. Financial support from IFM2 is gratefully appreciated.

†Address correspondance to Lars Stentoft, HEC Montreal, 3000 chemin de la Cote-Sainte-Catherine, Montreal (Quebec) Canada H3T 2A7, or e-mail: lars.stentoft@hec.ca. Phone: (+1) 514 340 6671. Fax: (+1) 514 340 5632.

1 Introduction

In the seminal paper by Black & Scholes (1973) a closed form solution for the price of a European option is derived. Since then, the Black-Scholes formula has been celebrated as one of the major successes of modern financial economics, although empirical analysis has pointed towards several systematic pricing errors when compared to observed option prices. To be specific, numerous studies have documented so-called smiles in the implied volatility as a function of moneyness as well as a tendency for the constant volatility model to underprice in particular short term out of the money options. In response to these "empirical regularities", a number of alternative models have been developed. In particular, the assumptions underlying the Black-Scholes model have been widely criticized, and much effort has been put into extending the valuation framework. Apart from the assumption of continuous trading, a crucial assumption in the Black-Scholes model is that of constant volatility and lognormality. However, the constant volatility lognormal model fails to accommodate a number of important features of asset return series, the most important of which are leptokurtosis and the volatility clustering phenomenon (see Bollerslev, Engle & Nelson (1994)).

In the option pricing literature, some of the earliest extensions to the Black-Scholes model are the continuous time stochastic volatility, or SV, models of Hull & White (1987), Wiggins (1987), Scott (1987), Stein & Stein (1991), and Heston (1993). More recent studies include, to name a few, Bakshi, Cao & Chen (1997) and Bates (2000). In these papers, the volatility is modelled as a separate stochastic process allowing for a large degree of flexibility in the specification. A particularly appealing feature of these models is that under certain assumptions elegant solutions can be derived for the price of European options. In some cases, the pricing formulas are approximately closed form solutions of the same general type as the Black-Scholes formula with integrals of the stochastic volatility. More generally, option prices may be available in semiclosed form through Fourier inversion of the underlying characteristic function (see e.g. Heston (1993), Bakshi et al. (1997), and Bates (2000)). However, in real applications a problem with the continuous time SV models is that volatility is unobservable, and hence estimation of these models is rather complicated. Examples of feasible estimation procedures include the Efficient Method of Moments, or EMM, method of Gallant & Tauchen (1996) and the Markov Chain Monte Carlo, or MCMC, method proposed by Jacquier, Polson & Rossi (1994) and Jacquier, Polson & Rossi (2004) (see also Johannes & Polson (2010)). Moreover, when pricing claims for which numerical procedures are required, as it is the case with e.g. American options, future values of the unobservable volatility are needed. This variable is latent and hence potentially complicated to predict, and thus such procedures may require e.g. the full reprojection machinery associated with the EMM procedure (see Gallant & Tauchen (1998)).

In the time series literature, several competing discrete time asset return models have been developed which can take account of the empirical features observed for financial returns. A large number of these extensions fall within the framework of autoregressive conditional heteroskedastic, or ARCH, processes suggested by Engle (1982) and the generalized ARCH, or GARCH, models introduced by Bollerslev (1986). A particularly appealing feature of these models is that data is readily available for estimation and this can be done with simple maximum likelihood procedures. These models have been successfully applied to financial data such as stock return data as seen from Bollerslev, Chou & Kroner's (1992) survey article. However, when it comes to option pricing a drawback with many of these models is that closed or approximately closed form solutions do not exist, although see Heston & Nandi (2000) for an exception. Thus, although the appropriate dynamics were derived in Duan (1995), for most of the existing models numerical methods have to be used for the actual pricing. For example, Monte Carlo simulation methods have been used to price options on the Standard and Poor's 500 Index using various GARCH specifications in Christoffersen & Jacobs (2004) and Hsieh & Ritchken (2005). Also, Bollerslev & Mikkelsen (1996) and Bollerslev & Mikkelsen (1999) have successfully used the GARCH option pricing framework together with fractionally integrated GARCH processes to price long-term European style equity anticipation securities (LEAPS) on this particular index.

However, in spite of the recent progress in terms of available estimation techniques for continuous time SV models and numerical methods for option pricing in the discrete time GARCH models, there is still a gap between the two strands of the literature. In particular, it is probably fair to say that in applications it is far from the standard to report estimation results for both discrete time and continuous time return models. Likewise, when pricing options with a continuous time SV model the results are usually only compared to other models formulated in continuous time, and when considering discrete time GARCH models the results are usually only compared to other models formulated in discrete time. Thus, the important question of which model is preferable remains open. In this paper, we bridge the gap between the discrete time approach and the continuous time approach using the Augmented GARCH model of Duan (1997) as the basic framework. To be specific, we choose a number of GARCH models within this framework, for which well known stochastic volatility models are obtained in the limit. In doing so, it becomes possible to compare directly the option prices calculated with discrete time models to those calculated from continuous time models.

The contribution of this paper is threefold: First of all, we provide estimation results for the GARCH models as well as for the diffusion limit SV models using simple maximum likelihood techniques. The estimates for the SV models are obtained by appropriately re-parametrizing the discrete time specifications and hence the two types of models are implemented in an internally consistent manner. Secondly, we compare the option price estimates from the discrete time models to their continuous time counterparts through a Monte Carlo study. We also examine the convergence in terms of time discretization and the effect of increasing the number of potential early exercise times. The results show that there are in fact differences between the discrete time GARCH models and their continuous time counterparts, though the discrete time GARCH prices converge quickly to the continuous time SV values. In terms of allowing for multiple intraday early exercises, the results show that this is potentially important for in the money options. Thirdly, the paper contains a large scale empirical analysis using options on a number of individual stocks and a stock index. The analysis shows that in actual applications of the models, the differences in overall performance are small. In general, this holds through time and across maturity and moneyness. The exception to this is for the in the money put options on the individual stocks, for which the continuous time specifications outperform the models formulated in discrete time. For this subsample of options, we also show that allowing for multiple intraday early exercises improves on the pricing performance.

The rest of the paper is organized as follows: In Section 2 the framework is introduced and estimation results are provided. In Section 3 the risk neutral dynamics are derived and we describe how options can be priced using simulation methods. In Section 4 the results of an extensive Monte Carlo study of the properties of the option pricing models are reported, and in Section 5 the option pricing models are taken to the data. Section 6 concludes.

2 Asset return model

In this paper, we consider a discrete time economy with the price of an asset denoted S_t and the dividends of that asset denoted d_t. We assume that the continuously compounded return process, $R_t = \ln(S_t + d_t) - \ln S_{t-1}$, can be modelled using the GARCH framework. The specific parametrization we use is

$$R_t = r + \lambda\sqrt{h_t} - \frac{1}{2}h_t + \sqrt{h_t}\varepsilon_t, \text{ with} \tag{1}$$

$$h_t = \omega + \alpha h_{t-1}(\varepsilon_{t-1} + \gamma)^2 + \beta h_{t-1}, \tag{2}$$

where $\varepsilon_t | \mathcal{F}_{t-1} \sim N(0,1)$, with $\mathcal{F}_{t-1}$ denoting the information set containing all information up to and including time $t-1$. It follows from lognormality that one plus the conditional expected rate of return equals $\exp\left(r + \lambda\sqrt{h_t}\right)$, and hence λ in (1) is readily interpreted as the unit risk premium when r is the continuously compounded risk-free rate of return. The model corresponds to the NGARCH model proposed by Engle & Ng (1993), which allows for the well known leverage effect through the parameter γ, and if $\gamma < 0$ such an effect is said to be found. It is clear that this model nests the ordinary GARCH specification which obtains when $\gamma = 0$.

2.1 Diffusion limits for the NGARCH model

The NGARCH model as well as the GARCH model are special cases of the Augmented GARCH framework of Duan (1997). This framework has been shown to contain as its limit several of the bivariate diffusion processes that are used as building blocks in various stochastic volatility models, see also Nelson (1990) for similar results for the GARCH model. In this section, we explain how this framework can be used to obtain the diffusion limits of the NGARCH model. We also discuss how the results allow for straightforward estimation of the parameters of the resulting continuous time models using simple maximum likelihood techniques.

2.1.1 The diffusion limit of Duan (1997)

In order to study the diffusion limits, we follow Duan (1997) and rewrite the discrete time daily NGARCH model from above as

$$R_{ks}^{(n)} = \left(r + \lambda\sqrt{h_{ks}^{(n)}} - \frac{1}{2}h_{ks}^{(n)}\right)s + \sqrt{h_{ks}^{(n)}}\varepsilon_k\sqrt{s}, \text{ with} \tag{3}$$

$$h_{(k+1)s}^{(n)} - h_{ks}^{(n)} = \omega s + \left(\beta + \alpha\left(1+\gamma^2\right) - 1\right)h_{ks}^{(n)}s + \alpha h_{ks}^{(n)}\left(\left(\varepsilon_k + \gamma\right)^2 - \left(1+\gamma^2\right)\right)\sqrt{s}, \tag{4}$$

where $s = 1/n$ is the length of the daily subintervals and $\varepsilon_k \sim N(0,1)$. Note that the specification in (1) and (2) obtains when $n = s = 1$. The limit is now considered as the length of the daily subintervals, s, tends to zero. From Duan (1997, Theorem 3) it can be shown that the limiting diffusion model of the system in (3) and (4) is characterized by

$$d\ln S_t = \left(r + \lambda\sqrt{h_t} - \frac{1}{2}h_t\right)dt + \sqrt{h_t}dW_{1t}, \text{ and} \tag{5}$$

$$dh_t = \omega dt + \left(\beta + \alpha\left(1+\gamma^2\right) - 1\right)h_t dt + 2\gamma\alpha h_t dW_{1t} + \sqrt{2}\alpha h_t dW_{2t}, \tag{6}$$

where W_{1t} and W_{2t} are two independent Wiener processes. This model corresponds to the bivariate diffusion model used in Hull & White (1987). In the following, we will refer to it as the C-NGARCH specification. The C-GARCH model, which obtains when $\gamma = 0$, has

$$dh_t = \omega dt + (\beta + \alpha - 1)\, h_t dt + \sqrt{2}\alpha h_t dW_{2t}. \tag{7}$$

This model corresponds to the special case of the Hull & White (1987) model where the volatility process is independent of the return process.

It should be noted that the diffusion limits of Nelson (1990) and Duan (1997) may not be unique. For example, Corradi (2000) shows that for some parameterizations one obtains a degenerate limit for the GARCH model. Moreover, Heston & Nandi (2000) shows that for their affine NGARCH model the limiting behavior is very different from that of the classical GARCH model as the same process drives both the spot asset and variance dynamics. However, since the diffusion limits depend on the parameterization used to obtain the limiting results, we are to a certain degree free to choose which limits to consider. The benefit of considering the results of Duan (1997) is that the derived limits correspond to models which are very popular in the continuous time option pricing literature. Thus, by selecting the limits carefully we can study in a simple way the relationship between the GARCH process and these well known limits when it comes to option pricing.[1] Moreover, though the derived limits in Corradi (2000) and Heston & Nandi (2000) allow options to be valued solely using the hedging arguments of Black & Scholes (1973) and Merton (1973), they are much less appealing from an empirical perspective. For example, if we were to use the limiting results of Corradi (2000), where the resulting diffusion is degenerate and hence does not allow for time varying volatility, the obtained results would differ greatly. However, this model is much less flexible and clearly not appropriate for financial asset return series.

[1] Ritchken & Trevor (1999) also consider the limits derived by Duan (1997). However, they report results for European options only and do not conduct an empirical exercise like we do.

2.1.2 Implementation and parameter estimation

The immediate benefit of working with the framework above is that the parameters of the diffusion model can be implied from the discretely observed data. In particular, the diffusion limit dynamics in (5) and (6) only depend on the parameters of the original discrete time system in (1) and (2). Thus, since the discrete time parameters are readily available with simple maximum likelihood estimation techniques, so are the implied parameters of the diffusion limit. However, an alternative procedure is to estimate the diffusion parameters directly. To do this, we first re-parametrize the diffusion limit of the NGARCH specifications as

$$d\ln S_t = \left(r + \lambda\sqrt{h_t} - \frac{1}{2}h_t\right)dt + \sqrt{h_t}dW_{1t}, \text{ and} \quad (8)$$

$$dh_t = \beta_0 dt + \beta_1 h_t dt + \beta_2 h_t dW_{1t} + \beta_3 h_t dW_{2t}, \quad (9)$$

where $\beta_0 = \omega$, $\beta_1 = \left(\beta + \alpha\left(1+\gamma^2\right) - 1\right)$, $\beta_2 = 2\gamma\alpha$, and $\beta_3 = \sqrt{2}\alpha$. Next, we rewrite the model in (1) and (2) in terms of these parameters to obtain

$$R_t = r + \lambda\sqrt{h_t} - \frac{1}{2}h_t + \sqrt{h_t}\varepsilon_t, \text{ with} \quad (10)$$

$$h_t = \beta_0 + \left(\beta_3/\sqrt{2}\right)h_{t-1}\left(\varepsilon_{t-1} + \left(\beta_2/\sqrt{2}\beta_3\right)\right)^2 + \left(1 + \beta_1 - \beta_3/\sqrt{2} + \beta_2^2/\sqrt{2^3}\beta_3\right)h_{t-1}. \quad (11)$$

The parameters in the model in (10) and (11) can also be estimated using simple maximum likelihood techniques, though compared to the parametrization in (1) and (2) the approach yields directly the implied continuous time parameters. Note that with this parametrization the GARCH model obtains when $\beta_2 = 0$.

A few comments are relevant with respect to the outlined estimation approach. First of all, our method relies on the availability of diffusion limits for the selected discrete time model. Though the Augmented GARCH framework of Duan (1997) contains many other discrete time models, it is possible that the diffusion limit of the discrete time process selected in a careful empirical analysis of the series to be analyzed is not known. Alternatively, one may wish to consider a continuous time model which does not correspond to the diffusion limit of a known discrete time model. In these situations one would have to use an alternative estimation procedure. Secondly, though the models in (1) and (2) and in (10) and (11) are estimated on the same data and using the same approach, there is no guarantee that they fit the data equally well. In particular, the model in (10) and (11) is highly nonlinear and this could lead to differences. Finally, the estimation procedure outlined above may not be fully efficient for the parameters of the diffusion limit as it does not allow for an additional shock.[2] Fully efficient estimates could be obtained with e.g. the EMM or MCMC methods, though these methods are computationally much more complex. However, since these methods are often implemented using daily data our approach may nevertheless give a good idea about the model performance.

2.2 Empirical results

In this section, we use the above framework to study a sample of financial assets. The next section introduces the data which has been previously studied in Stentoft (2005) and Stentoft (2008). We then provide estimation results for the discrete time GARCH and NGARCH models as well as for the diffusion limits of these models, both of which are based in simple maximum likelihood techniques.

2.2.1 Data

For the empirical work, we use data for General Motors (GM), International Business Machines (IBM), Merck and Company Inc. (MRK), as well as for the Standard and Poor's 100 Index (OEX). The reason for choosing these three stocks is that for the period under consideration options on them were the most traded in terms of actual trades as well as in terms of total volume. The reason for choosing the Standard and Poor's 100 index (OEX) is that this is the broadest index for which

[2] In this respect our proposed approach is similar to the one used in e.g. Fleming & Kirby (2003) for estimating discrete time stochastic volatility models using GARCH filters.

options are traded on the CBOE and it has been the focus of much research.

The return series for the individual stocks were obtained from the Center for Research in Security Prices (CRSP). We use return data beginning January 2, 1976, since this is as far back data on the individual stock returns and dividends are available to us on a daily basis. The continuously compounded return series in percentage terms for the Standard and Poor's 100 Index was calculated from the return index supplied by Datastream. Since our data on the corresponding options ends December 29, 1995, this date also marks the end of the sample which as a result contains 5055 daily observations.

Table 1 shows sample statistics and Figure 1 provides time series plots for the four return series. From the table it is seen that the returns are generally negatively skewed and leptokurtic, although for MRK the skewness is insignificantly different from zero. From the figure it is seen that the returns are clearly not independently and identically distributed through time. On the contrary, periods of low volatility are followed by high volatility periods and vice versa, a finding known as volatility clustering. The GARCH framework has been successfully applied to data with these characteristics, see e.g. Bollerslev et al.'s (1992) survey article.

2.2.2 Estimation results

Tables 2 to 5 report Quasi Maximum Likelihood (QML) estimation results for the model in (1) and (2). First of all, columns four and five in the tables show the estimation results for the GARCH models. Compared to the simpler CV model in columns two and three, the tables show that allowing for time varying volatility leads to large increases in the Log-Likelihood values. Furthermore, for all series both extra parameters, α and β, are estimated significantly different from zero, and the estimates are in line with what is usually found in the literature. In terms of serial correlation in the squared standardized residuals, the $Q^2(20)$ statistics show that for all but GM the null of no correlation cannot be rejected for this model. Thus, it seems that modelling volatility as a GARCH process goes quite a way in terms of eliminating the ARCH effects for IBM, MRK, and

OEX. Furthermore, for three of the four series the $Q(20)$ statistics are now insignificant at a one percent level.

Secondly, columns six and seven of the tables present estimation results for the NGARCH model. The tables show that in all cases adding the leverage parameter γ leads to large increases in the Log-Likelihood value. Furthermore, for all the return series the estimated value of γ is significantly different from zero and has the expected sign. In terms of the diagnostic tests, however, adding the asymmetry parameter does not change a lot except for OEX where the $Q(20)$ statistics is now also insignificant at a five percent level. The Schwarz Information Criteria, SIC, value is smaller for the asymmetric models than for the symmetric GARCH model. This indicates that asymmetries in the volatility specification are important features of the return data under consideration and that this type of model should be preferred.

Next, we consider the diffusion limits of the GARCH and NGARCH models above. In columns nine and ten and eleven and twelve, respectively, of Tables 2 to 5 we report QML estimation results for the model in (10) and (11), that is the model re-parametrized directly in terms of the diffusion limit parameters. Using this specification for estimation instead of simply implying the parameters allows for direct testing of the parameters of the diffusion limits.[3] From the tables we first of all note that overall the estimated parameters seem very reasonable. In particular, they show that there is strong persistence in the shocks to the variance process in the diffusion limits. Secondly, when allowed for asymmetries are found to be highly significant, and the parameter estimates imply a very large and negative correlation between the return and variance processes for all four series.

Finally, when comparing the results for the discrete time and continuous time models in Tables 2 to 5 it is seen that the highly nonlinear parameter transformations imposed in the continuous time specification do not lead to any decrease in the likelihood values. Moreover, we obtain the same parameters if we imply these from the discrete time model in (1) and (2) instead of estimating them

[3] If the parameters were implied instead, the delta method would have to be used to obtain the appropriate standard errors.

directly using (10) and (11).[4] On the other hand, since the estimation procedure does not allow explicitly for an additional shock in the diffusion model, it is not surprising that the statistical fit of the continuous time models is virtually identical to that of the discrete time models. However, a more detailed econometric analysis of the results is beyond the scope of the present paper, and at this time we refrain from commenting further on the parameters. Our metric is after all one of option pricing performance which we turn towards now.

3 Risk neutral dynamics

The discrete time models used in this paper can all be written in the following general form:

$$R_t = r + \lambda\sqrt{h_t} - \frac{1}{2}h_t + \sqrt{h_t}\varepsilon_t, \text{ with} \tag{12}$$

$$h_t = g\left(\theta, h_s, \varepsilon_s; s \leq t-1\right), \tag{13}$$

where $\varepsilon_t | \mathcal{F}_{t-1} \sim N(0,1)$ under measure $\mathcal{P}$. In (13), θ denotes the set of parameters used to specify the variance dynamics. For example, for the NGARCH specification in (2) we have $\theta = \{\omega, \alpha, \beta, \gamma\}$. To use this model for option pricing purposes we use the Locally Risk-Neutral Valuation Relationship (LRNVR) derived in Duan (1995) which can be shown to hold under some familiar assumptions on preferences and assumed conditional lognormality. Invoking it, the dynamics to by used for option pricing are easily shown to be given by

$$R_t = r - \frac{1}{2}h_t + \sqrt{h_t}\varepsilon_t^*, \text{ with} \tag{14}$$

$$h_t = g\left(\theta, h_s, \varepsilon_s^* - \lambda; s \leq t-1\right), \tag{15}$$

where $\varepsilon_t^* | \mathcal{F}_{t-1} \sim N(0,1)$ under measure $\mathcal{Q}$. Thus, the risk neutral dynamics depend only on the parameters in the original volatility specification θ and the unit risk premium λ. Since all of the

[4] These findings are robust to using different starting values as well as to using different optimization algorithms.

necessary parameters can be estimated from asset returns obtaining the risk neutral dynamics is straightforward.

Equations (14) and (15) show that the dynamics remain Gaussian though with a shifted mean. The shift in mean corresponds to what is needed to compensate investors for holding the risky assets. Note that, while we may be lured into believing that we have successfully eliminated all preference related parameters this is not the case. However, the LRNVR is sufficient to reduce the preference considerations to the constant unit risk premium λ present in the variance equation.

3.1 Option pricing with the limiting diffusion

As it is the case under the data generating process, the Augmented GARCH process under the risk-neutralized pricing measure can be shown to converge to a bivariate diffusion system. This was shown in Duan (1996, Theorem 2), and this general theorem can be used to derive a system corresponding to the risk-neutralized version of the GARCH variance specifications used above. Alternatively, it is possible to show that the risk-neutral dynamics can be derived directly from the corresponding diffusion limits under the data generating process.

To fix ideas, assume that the diffusion limit of the Augmented GARCH system in (12) and (13) has been derived under measure $\mathcal{P}$. With a slight abuse of notation we specify this as

$$d\ln S_t = \left(r + \lambda\sqrt{h_t} - \frac{1}{2}h_t\right)dt + \sqrt{h_t}dW_{1t}, \text{ and} \tag{16}$$

$$dh_t = g\left(\theta, h_t, dW_{1t}, dW_{2t}\right), \tag{17}$$

where W_{1t} and W_{2t} are the two independent Wiener processes from above, and where we again let θ denote the set of parameters used to specify the variance dynamics. It then follows from Duan (1996) that the diffusion limit under the corresponding risk-neutralized pricing measure, $\mathcal{Q}$, becomes

$$d\ln S_t = \left(r - \frac{1}{2}h_t\right)dt + \sqrt{h_t}dW_{1t}^*, \text{ and} \tag{18}$$

$$dh_t = g\left(\theta, h_t, dW_{1t}^* - \lambda dt, dW_{2t}^*\right), \tag{19}$$

where W_{1t}^* and W_{2t}^* are two independent Wiener processes. Once again, it is straightforward to obtain these risk neutral dynamics since all of the necessary parameters can be estimated from asset returns.

We note that, as it was the case in the discrete time system in (14) and (15), the dynamics remain Gaussian but with a shifted mean in one of the innovation terms in the variance process. Thus, on the one hand the result in (18) and (19) confirms the result of the original paper by Hull & White (1987). In particular, it implies that when the two innovations are independent the premium for volatility risk is zero. On the other hand, it extends the results and provides the risk neutral dynamics for the case of correlated innovations. For further discussion of these issues see Duan (1996).

3.2 Implementation of the GARCH option pricing model using simulation

Although the pricing system in (14) and (15), or for that sake the system in (18) and (19), is completely self-contained, an actual application to even the simple European option is difficult because of the lack of a closed form expression for the value of the underlying asset at maturity of the option. However, it is immediately clear that using the system in (14) and (15), respectively that in (18) and (19), a large number of paths of the risk-neutralized asset prices can be generated, possibly by using a type of discretization scheme. From this sample of paths, an estimate of the European option value can be obtained as a simple average of the discounted pathwise final payoffs. For option pricing purposes, this method has been used at least since Boyle (1977).

For the American option things are not as simple since an optimal exercise strategy has to be determined simultaneously. However, the work by e.g. Carriere (1996), Longstaff & Schwartz (2001), and Tsitsiklis & Van Roy (2001) has shown how this can be done using a simulation approach. By now these methods have become standard tools in financial economics. The most important of these contributions in terms of their use is the Least Squares Monte Carlo (LSM) method of Longstaff & Schwartz (2001). In a GARCH context the LSM method was used successfully in Stentoft (2005) and Stentoft (2008). The only requirement for the method is that we are able to generate simulated paths from the appropriate risk neutral system. It can therefore be equally well used to price American options for the discrete time case described in (14) and (15) and for the continuous time case described in (18) and (19). Thus, in the present paper we use this particular algorithm to price American options in both of these cases.

The LSM method for pricing American style options proceeds as follows: First of all, given the full sample of random paths, the pricing step is initiated at the maturity date of the option. At this time, it is possible to decide along each path if the option should be exercised since the future value trivially equals zero. Hence, the pathwise payoffs may be easily determined at maturity. Next, working backward through time a cross-sectional regression is performed at the first point in time where early exercise is to be considered. In the regression the future pathwise payoffs in the simulation are regressed on transformations of the current pathwise asset prices and volatility levels. The fitted values from this regression are then used as estimates of the pathwise conditional expected values of holding the option for one more period. The decision of whether to exercise or not along each path can now be made by comparing the estimated conditional expected value of continuing to hold the option to the value of immediate exercise. In particular, if immediate exercise yields superior payoff, this is the optimal choice along this particular path. Once the decision has been recorded for each path, we can move back through time to the previous early exercise point and perform a new cross-sectional regression with the appropriate pathwise payoffs based on the previously determined choices. Finally, with the optimal early exercise strategies along each path an estimate of the American option value can be obtained as a simple average of the discounted pathwise payoff, as it is the case for the European option.

4 Option pricing properties

In this section, we conduct a Monte Carlo study which compares the estimated American and European option prices obtained with the discrete time models with what is obtained with the corresponding diffusion limits. We first of all report results for the discrete time GARCH and NGARCH specifications as well as for their corresponding diffusion limits which corresponds to the bivariate diffusion model in Hull & White (1987). While the latter results are interesting on their own, by carefully choosing the particular bivariate diffusion models the study allows us to compare the two approaches directly and in a consistent manner. Thus, this section extends the results in Stentoft (2005) to the continuous time framework. Next, because of the careful choice it is possible to gauge how well the family of discrete time GARCH models actually approximates the continuous time SV models in terms of option pricing by decreasing the size of the steps in the GARCH and NGARCH specifications. Finally, using either of the specifications it is possible to examine the convergence of the price estimates as the number of early exercise points is increased.

4.1 Data and parameter values

To illustrate the pricing properties we consider a set of artificial put options with strike prices ranging from deep out of the money, a moneyness equal to 0.90, to deep in the money, a moneyness of 1.10, and with ultra short (7 trading days), short (21 trading days), middle (63 trading days) and long maturities (126 trading days). We take a year to be 252 trading days. To a large extent this collection of options covers what is actually observed for traded options empirically. In line with what is found empirically, we set the interest rate equal to 6% on an annual basis, and for the time being we assume that dividend payments are zero.

For the dynamics we choose parameter values for ω, α, β, and γ which are empirically plausible.

To be specific, we consider values which are close to the actual averages of the estimated parameters from the return series from 1976 through 1995. The values are given in Table 6. Thus, we set $\beta = 0.92$ and $\alpha = 0.06$ in the GARCH specification, and $\beta = 0.92$, $\alpha = 0.048$, and $\gamma = -0.5$ in the NGARCH model. These values yield a persistence of 0.98 for both models. Moreover we set $\omega = 4.96 \times 10^{-6}$ which implies an annualized unconditional volatility of 25%. Finally, we fix $\lambda = 0.05$ and to start up the simulations we set $E[h]$ equal to the unconditional level of the variance. In Table 7 we show the implied parameter values for the diffusion limit models.

4.2 Option pricing properties

We now provide option prices for the different models using the discrete time and the implied continuous time specifications with the parameter values in Table 6 and 7, respectively. The reported results are averages of 100 calculated estimates using different seeds in the random number generator. The standard errors of these 100 estimates are reported in parentheses below the corresponding price estimate. In each simulation $100,000$ paths are used. For the American price estimates in Table 8, powers of and cross products between the asset price and the level of the volatility of total order less than or equal to three are used in the cross-sectional regressions, which are used to estimate the continuation value for the in the money paths when early exercise is considered. For the time being, we assume that early exercise is only considered at the end of each trading day, an assumption which is relaxed in Section 4.3.

4.2.1 American option pricing results using GARCH specifications

Column four in Table 8 reports option prices calculated using the GARCH specification, and column five shows the relative bias which would arise if one was to use a constant volatility specification like in the model of Black & Scholes (1973) for option pricing instead. Thus, the bias indicates the mispricing by the CV model which would be observed if the true model is in fact the GARCH model. From this column it is clear that out of the money short maturity options would be particularly underpriced by the CV model, whereas at the money options would be overpriced. Although these effects persist as the maturity increases, the relative bias and particularly the underpricing of the out of the money options become less and less pronounced.

Column six in Table 8 shows the NGARCH option prices and column seven the relative bias arising from using the CV model to price the options. From these columns it is clear that the underpricing of out of the money options by the CV model is even more pronounced than with the GARCH model. Furthermore, once asymmetries are introduced in the volatility model, mispricings are present for even the longest maturities considered. In fact, the underpricing of the CV model relative to the NGARCH model for long term out of the money put options is 36.56%.

4.2.2 American option pricing results using GARCH diffusion limit specifications

The last four columns in Table 8 present the results obtained when the diffusion limits of the GARCH and NGARCH models are used. In order to simulate the stock and variance paths an Euler discretization of the appropriate processes with 1024 intraday steps is used. Note that with this fine discretization the computational time increases by a factor of 16. Compared to columns four through seven, with the discrete time prices, we see that the same pricing pattern results when using the diffusion limits of the GARCH models as with the actual GARCH models. In particular, the underpricing of out of the money options in the CV model remains very pronounced for the NGARCH diffusion limit.

More importantly though, Table 8 shows that there are differences between the price estimates obtained with the discrete time models and those from the implied diffusion models. In particular, this is the case for the NGARCH specification, and the table shows that for almost all the options the diffusion model yields higher price estimates. The differences increase with maturity, and for the options with the longest maturity the average relative difference is 1.4%. For the GARCH diffusion limit differences also occur reflecting the lack of correlation between the return and variance processes in the diffusion limit.

4.2.3 Pricing results for the European options and early exercise values

Though our main focus is on pricing American options, for completeness we report the corresponding results for the European options in Table 9. When comparing the values in this table to those in Table 8 it is seen that the results are very similar. In particular, for the discrete time models Table 9 shows that out of the money short maturity options would be particularly underpriced by the CV model, and once asymmetries are introduced in the volatility model mispricings are present for even the longest maturities considered. Moreover, the table also shows that there are differences between the European price estimates obtained with the discrete time models and those from the implied diffusion models. Again these differences are similar to what was found for the American options in Table 8.

By comparing the actual price estimates in Tables 8 and 9 one can calculate the estimated early exercise value as the difference between the American and the European price. The results are shown in Table 10, and once again it is seen that compared to the models with time varying volatility the CV model underestimates these early exercise values. Moreover, the table shows that there are also differences between the estimates obtained with the discrete time models and those from the implied diffusion models. In particular, the latter models on average produce the largest estimated early exercise values. Again, the relative differences increase with maturity and for the options with the longest maturity the average relative difference is 9.6%. Finally, it is worth noting that for this particular sample of options the average early exercise values are around 30 cents and in relative terms may be as large as 12%. Thus, early exercise is clearly important and should not be neglected for the models considered here.

4.3 Convergence properties

Above we showed that there are indeed differences when it comes to option prices calculated using the discrete time specifications and the implied continuous time limits. We now analyze the option prices as the GARCH specification converges to the diffusion limit. While the convergence properties have been studied for the asset dynamics as such and more recently also for the European price estimates, see e.g. Duan, Wang & Zou (2009), this is not the case for the American option pricing problem. First of all, we study how the price estimates based on discrete time models converge to estimates based on the continuous time specification as the number of steps per day is increased while keeping the number of daily early exercise points fixed at one. Secondly, we study the convergence of the option price as the number of early exercise points is increased.

4.3.1 Convergence to continuous time diffusion

As the benchmark continuous time specification, we take the values obtained using an Euler discretization with 1024 intraday steps available in Table 8. We compare these prices to what is obtained with the discrete time model when the number of steps per day is increased and calculate the pricing errors as the difference between the discrete time model price and the continuous time benchmark price. In Figure 2 we plot the pricing errors as a function of the strike price and the logarithm of the number of intraday steps for the American option prices with one early exercise day obtained with the NGARCH model. Each plot is for one particular time to maturity: From left to right and top to bottom the plots are for long term, or LT, options with $T = 126$ days to maturity, for medium term, or MT, options with $T = 63$ days to maturity, for short term, or ST, options with $T = 21$ days to maturity, and for ultra short term, or UST, options with $T = 7$ days to maturity.

Figure 2 shows that the convergence in all the plots is monotone and very smooth. Moreover, it happens quite quickly. In fact, considered across all maturities and strike prices the average error is less than one cent when as little as 32 intraday steps are used in the discrete time approximation. For the GARCH model, the convergence pattern is similar and we therefore refrain from reporting results, and for the CV model the simulation is exact and therefore there is no issue of convergence for this model. Compared to the convergence of the European price estimates shown in Figure 3 the pattern is very similar.[5] This also holds when considering American options where multiple early exercises each day are allowed (the results are available from the author upon request). However, it should be noted that as the number of early exercise times is increased the initial differences increase as well.

4.3.2 Convergence to continuous time early exercise

In the previous sections we considered options with a fixed finite number of early exercise points. However, American options can be exercised at any point prior to maturity. The framework we use here allows us to analyze the convergence properties of our price estimates as the number of early exercises is increased. Note that as we keep the number of paths and the number of regressors constant convergence is to the corresponding approximation of the American option price (see also the results in Stentoft (2004)). As the benchmark American continuous time specification we take the values obtained using an Euler discretization with early exercise considered at each of the steps. Due to the computational complexity, we limit attention to the shortest term options with $T = 21$ and $T = 7$ days to expiration and with a maximum of 64 and 256 intraday steps, respectively.

The first issue we examine is the convergence of the Bermudan style option estimates to the American ones. To focus on this aspect we compute option prices with increasing number of early exercise points while keeping the overall Euler discretization in the continuous time model at the maximum number of intraday steps. The results are shown in the top row in Figure 4, which plots the pricing errors for the NGARCH model, defined as the price of the Bermudan option minus the benchmark continuously exercisable American price, as a function of the strike price and the logarithm of the number of intraday steps. The plots show that there is indeed an effect of increasing the number of early exercise points. The effect is largest in dollar terms for options

[5] In Duan et al. (2009) bounds on the convergence rate are derived and these are shown to be up to an order near the square root of the length of the time interval. Using the simulated prices in our Monte Carlo study it is possible to estimate the convergence rates numerical. When doing so, we note that the derived bounds in Duan et al. (2009) for the European prices are satisfied. Moreover, when estimating the rate of convergence for the American option prices it is essentially the same as that obtained for the European prices.

which are in the money. However, note that even for these options the convergence occurs rapidly.

The second issue we examine is the convergence of the discrete time Bermudan style option value to the continuous time American option value. That is, in this case we examine the convergence when both the number of intraday steps and the number of intraday early exercises are increased. The result of this is shown in the bottom row in Figure 4 which plots the pricing errors for the NGARCH model, defined as the price of the discrete time Bermudan option minus the benchmark continuously exercisable American price, as a function of the strike price and the logarithm of the number of intraday steps. Note that this benchmark is the same as before, and therefore from these plots the combined effect can be gauged. The plots show that for the at the money options the convergence in terms of the number of intraday steps is most important, whereas for the in the money options it is the convergence in terms of intraday early exercises which is most important.

5 Empirical performance

In this section, we take our model to the data and price a large sample of American style options on the three individual stocks as well as on the S&P 100 Index. We start by describing the data and the methodology used. We then provide results on the overall performance of the discrete time and continuous time versions of the models, and we analyze this performance through time as well as across maturity and moneyness. Finally, we consider a subset of the data, the in the money options, and examine the benefit of allowing for multiple intraday early exercises for options of this type.

5.1 Data, estimation, and pricing methodology

The option data we use covers the period from 1991 through 1995 and contains weekly observations, which are sampled on Wednesdays or the closest day to Wednesday if this is a holiday. The reason for using weekly observations only is to trade-off having a relatively long time period against having a reasonable amount of computational work. We apply several filters to the data as detailed in Stentoft (2005). At each date of pricing t the models are re-estimated such that for pricing purposes all the historical observations available at time t are considered. Note that this would be computational demanding unless efficient estimation methods were available. However, this is not a problem here since simple maximum likelihood procedures can be used to obtain the parameters of both the discrete time and continuous time models.

In addition to allowing for easy estimation, the discrete time models are also natural candidates for filtering out the volatility process needed for option pricing purposes. Thus, with our framework the proposed simulation method can be easily implemented. In the pricing procedure, we use a total of $M = 100,000$ paths in the simulations, and for the continuous time diffusion limit models we use 32 discretizations per trading day to limit the amount of computational time. Increasing the number of steps to 64 did not substantially change the results. With this choice, the computational time for the continuous time models is roughly twice that of the discrete time models for a given number of early exercise times. Finally, when estimating the conditional expectations for the in the money paths in the pricing step, we use powers of and cross products between the asset price and the level of the volatility of order three or less in addition to a constant term in the LSM procedure for both the discrete time and continuous time models.

In the simulations, we make the following three assumptions about the effect of dividend payments: First, we assume that only cash dividend payments are important for our purpose. This assumption is reasonable since exchange traded options, in general, are protected against other forms of dividends like, say stock splits. Secondly, we assume that both the ex-dividend day and the size of the dividends are known in advance. Though this is not strictly correct, dividends are paid regularly with fairly stable amounts throughout the period we consider. Thirdly, we assume that the effect of a cash dividend payment fully spills over on the asset price. Thus, if day t is an ex-dividend day and the simulated risk-neutral continuously compounded return is R_t, the end of day asset price is calculated from the price at the previous day as

$$S_t = S_{t-1} * \exp(R_t) - d_{t-1}. \tag{20}$$

We note that treating cash dividend payments as known both in size and timing and letting the payment fully spill over on the asset price is standard procedure.

5.2 Overall performance of the GARCH option pricing model

We price options using both the discrete time models and the implied continuous time models assuming for the present that the options can be exercised only once a day. In the literature, a number of different metrics have been used to gauge the performance of alternate option pricing models (see e.g. Bollerslev & Mikkelsen (1999)). We choose to report results on two of these, the bias, $BIAS \equiv K^{-1}\sum_{k=1}^{K}\left(\bar{P}_k - P_k\right)$, and the root mean squared error, $RMSE \equiv \sqrt{K^{-1}\sum_{k=1}^{K}\left(\bar{P}_k - P_k\right)^2}$, where P_k denotes the k'th observed price and $\bar{P}_k$ denotes the k'th price estimate. We also report results in terms of the implied standard errors, or ISD, using the same metrics. The ISDs are backed out from the Binomial Model with daily early exercise and corrects for maturity and moneyness effects through the nonlinear transformation of the dollar price. Tables 11 and 12 provide results for each of these metrics for the dollar errors and ISD errors, respectively.

The first thing to note from Table 11 is that when the dollar pricing errors from the discrete time GARCH models are compared to the CV models, the time varying volatility models are the preferred ones. This holds for all assets irrespective of wether put or call options are considered and irrespective of the type of metric used. The table also shows that the same conclusions hold for the continuous time models. Next, Table 12 shows that the same conclusion holds when ISD errors are considered. Finally, Tables 11 and 12 also show that within both types of framework asymmetries are important.

When comparing the discrete and continuous time models, Tables 11 and 12 show that the actual pricing errors are very close. In particular, for the dollar errors the maximum difference in the BIAS is 2.3 cents, and for the ISD errors the difference between the discrete time and continuous time BIAS is in all cases less then 25 basis points. The largest differences are in terms of RMSE, where the continuous time specifications outperform the discrete time models by about 10% for put and call options on IBM in terms of dollar losses, and where the discrete time GARCH specification outperforms the diffusion limit by about 11% for put options on OEX. Thus, it is not immediately clear if the increased computational complexity of the continuous time models, which is roughly twice that of the discrete time models, is reflected in added precision.

5.3 Maturity and moneyness effects for the options

We now analyze the performance of the models through time as well as across maturity and moneyness. In all cases, we report results on the ISD BIAS and we group all the individual stocks together (detailed results are available from the author upon request). Through time, we split the sample by year for a total of 5 subsamples. In terms of maturity, we split the sample in ultra short term options, or UST with $T < 11$, short term, ST with $11 \leq T < 21$, medium term, or MT with $21 \leq T < 63$, long term, or LT with $63 \leq T < 126$, and ultra long term, or ULT with $126 \leq T$, where T is the days to maturity. Finally, in terms of moneyness we split the sample in deep in the money options, or DITM which are 6% or more in the money, in the money options, or ITM which are between 2% and 6% in the money, at the money options, or ATM which are between 2% in and 2% out of the money, out of the money options, or OTM which are between 2% and 6% out of the money, and deep out of the money options, or DOTM which are 6% or more out of the money.

The results across the three dimensions are shown in Figures 5, 6, and 7, respectively. The figures first of all show that across most dimensions the time varying models clearly outperform the CV benchmark model. In particular, this is the case both in terms of the level of the errors and in terms of the variation in performance across each dimension. However, the figures also show that once time varying features are included whether the models are formulated in discrete time or in continuous time appears to be of second order importance. Thus, the figures show that the close overall performance seen in Tables 11 and 12 generally holds across these different dimensions.

While Figures 5, 6, and 7 illustrate the close overall performance, they do show differences in model performance across the three dimensions. For example, for the index put options the C-GARCH model performs somewhat worse than the three other models with time varying volatility for the UST category and in 1991. However, the relative differences are in fact larger for put options on individual stocks, and the largest differences are found when considering the moneyness dimension in Figure 7. In particular, for the deep in the money category of options the error for the C-NGARCH model is as little as -0.29% which is less than half that of the NGARCH model with an error of 0.77%.[6] The top left plot in Figure 7 shows that because of this difference in size and in sign, the C-NGARCH model virtually eliminates the variation in performance across the moneyness dimension for put options on individual stocks.

5.4 Multiple intraday early exercises

In Section 4 we analyzed the convergence properties of the price estimates when the number of early exercise times increases and showed that this is potentially important for options which are in the money. Moreover, the previous section showed that for this particular category of options the largest differences are found empirically between the discrete time and continuous time models. We therefore examine the potential benefit of allowing for multiple intraday exercises for this subsample of options, the individual stock put options which are in the money or deep in the money.

In Table 13 we provide details on the ISD errors for this subset of options when only one intraday early exercise is allowed for. Thus, these results correspond to what is reported in e.g. Figure 7 and the table confirms that for this subsample of options the diffusion limit models outperform the discrete time models and produces smaller pricing errors. In Table 14 results are shown for

[6] The exact values for the in the money and deep in the money options can also be found in Table 13.

the diffusion limit model when one and when multiple intraday early exercises are allowed for. The table shows that allowing for multiple intraday early exercises does affect the option pricing errors for this sample of options. In particular, the table shows that allowing for multiple daily early exercises further decrease the pricing errors for the C-NGARCH model, which was the best performing model for the DITM options. For this model the BIAS is very close to zero and the RMSE decreases by 4.2% when looking at Panel D.

6 Conclusion

This paper uses the Augmented GARCH framework of Duan (1997) to bridge the gap between the discrete time GARCH models and the continuous time SV models and uses these for American option pricing. We first of all provide estimation results for the GARCH models as well as for the diffusion limit SV models using simple maximum likelihood techniques. The estimates for the SV models are obtained by appropriately re-parametrizing the discrete time specifications and hence the two types of models can be implemented in an internally consistent manner.

We then perform a Monte Carlo study to examine the potential differences between the models in terms of option prices, and we study the convergence of the discrete time option prices to their implied continuous time values. The results show that there are in fact differences between the discrete time GARCH models and their continuous time SV counterparts. However, the differences are relatively small and the option price estimates from the discrete time models converge smoothly and quite quickly to those obtained with continuous time models. When multiple intraday early exercises are allowed for, the results show that this is most important for the in the money options though convergence to the continuous time American option value happens quickly.

Finally, a large scale empirical analysis is performed comparing the estimated prices based on discrete time models to the corresponding continuous time models. The results show that, while the differences in overall performance are small, for in the money options the continuous time SV models generally perform better than the discrete time GARCH specifications. Moreover, for this subsample of data there are potential gains in performance from allowing for multiple intraday early exercises. However, as the continuous time models are computationally more demanding in terms of option pricing, in particular when multiple intraday early exercises are considered, it may be argued that the discrete time GARCH models provide a very reasonable alternative.

Our approach and results constitute a first step towards bridging the gap between the discrete time and the continuous time models used for option pricing, and there are interesting extensions for future research. First of all, while the method used in this paper provides for easy estimation of the parameters of the diffusions, the method may not yield efficient estimates. Fully efficient estimates could be obtained with e.g. the EMM or MCMC methods, which exploit the continuous time structure of the model. Though these methods are computationally much more complex, it is possible that such estimates would lead to better performance of the continuous time models in terms of option pricing. Secondly, in our empirical application we do not consider option data when estimating the models and important information could therefore be neglected. Though incorporating option data is inherently difficult for American style derivatives, it is nevertheless an important issue to consider.

References

Bakshi, G., Cao, C. & Chen, Z. (1997), 'Empirical Performance of Alternative Option Pricing Models', *Journal of Finance* **52**(5), 2003–2049.

Bates, D. S. (2000), 'Post-'87 Crash Fears in the SP 500 Futures Option Market', *Journal of Econometrics* **94**, 181–238.

Black, F. & Scholes, M. (1973), 'The Pricing of Options and Corporate Liabilities', *Journal of Political Economy* **81**, 637–654.

Bollerslev, T. (1986), 'Generalized Autoregressive Conditional Heteroskedasticity', *Journal of Econometrics* **31**, 307–327.

Bollerslev, T., Chou, R. Y. & Kroner, K. F. (1992), 'ARCH Modelling in Finance', *Journal of Econometrics* **52**, 5–59.

Bollerslev, T., Engle, R. F. & Nelson, D. B. (1994), ARCH Models, *in* R. F. Engle & D. L. McFadden, eds, 'Handbook of Econometrics', Vol. 4, Elsevier, pp. 2959–3038.

Bollerslev, T. & Mikkelsen, H. O. (1996), 'Modelling and Pricing Long Memory in Stock Market Volatility', *Journal of Econometrics* **73**, 151–184.

Bollerslev, T. & Mikkelsen, H. O. (1999), 'Long-Term Equity Anticipation Securities and Stock Market Volatility Dynamics', *Journal of Econometrics* **92**, 75–99.

Boyle, P. P. (1977), 'Options: A Monte Carlo Approach', *Journal of Financial Economics* **4**, 323–338.

Carriere, J. F. (1996), 'Valuation of the Early-Exercise Price for Options using Simulations and Nonparametric Regression', *Insurance: Mathematics and Economics* **19**, 19–30.

Christoffersen, P. & Jacobs, K. (2004), 'Which GARCH Model for Option Valuation?', *Management Science* **50**(9), 1204–1221.

Corradi, V. (2000), 'Reconsidering the Continuous Time Limit of the GARCH(1,1) Process', *Journal of Econometrics* **96**, 145–153.

Duan, J.-C. (1995), 'The GARCH Option Pricing Model', *Mathematical Finance* **5**(1), 13–32.

Duan, J.-C. (1996), 'A Unified Theory of Option Pricing under Stochastic Volatility - from GARCH to Diffusion', *Manuscript* .

Duan, J.-C. (1997), 'Augmented GARCH(p,q) Process and its Diffusion Limit', *Journal of Econometrics* **79**, 97–127.

Duan, J.-C., Wang, Y. & Zou, J. (2009), 'Convergence Speed of GARCH Option Price to Diffusion Option Price', *International Journal of Theoretical and Applied Finance* **12**(3), 359–391.

Engle, R. F. (1982), 'Autoregressive Conditional Heteroscedasticity with Estimates of the Variance of United Kingdom Inflation', *Econometrica* **50**(4), 987–1007.

Engle, R. F. & Ng, V. K. (1993), 'Measuring and Testing the Impact of News on Volatility', *Journal of Finance* **48**(5), 1749–1778.

Fleming, J. & Kirby, C. (2003), 'A Closer Look at the Relation between GARCH and Stochastic Autoregressive Volatility', *Journal of Financial Econometrics* **1**(3), 365–419.

Gallant, A. R. & Tauchen, G. (1996), 'Which Moments to Match', *Econometric Theory* **12**, 657–681.

Gallant, R. & Tauchen, G. (1998), 'Reprojecting Partially Observed Systems With Application to Interest Rate Diffusions', *Journal of the American Statistical Association* **93**(441), 10–24.

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. & Nandi, S. (2000), 'A Closed-Form GARCH Option Valuation Model', *Review of Financial Studies* **13**(3), 585–625.

Hsieh, K. C. & Ritchken, P. (2005), 'An Empirical Comparison of GARCH Option Pricing Models', *Review of Derivatives Research* **8**, 129–150.

Hull, J. & White, A. (1987), 'The Pricing of Options on Assets with Stochastic Volatilities', *Journal of Finance* **42**(2), 281–300.

Jacquier, E., Polson, N. G. & Rossi, P. E. (1994), 'Baysian Analysis of Stochastic Volatility Models', *Journal of Business and Economic Statistics* **12**(4), 371–389.

Jacquier, E., Polson, N. G. & Rossi, P. E. (2004), 'Bayesian Analysis of Stochastic Volatility Models with Fat-tails and Correlated Errors', *Journal of Econometrics* **122**, 185–212.

Johannes, M. & Polson, N. (2010), MCMC Methods for Financial Econometrics, *in* Y. Ait-Sahalia & L. P. Hansen, eds, 'Handbook of Financial Econometrics', Vol. 2, Elsevier, pp. 1–72.

Longstaff, F. A. & Schwartz, E. S. (2001), 'Valuing American Options by Simulation: A Simple Least-Squares Approach', *Review of Financial Studies* **14**, 113–147.

Merton, R. C. (1973), 'Theory of Rational Option Pricing', *Bell Journal of Economics and Management Science* **4**, 141–183.

Nelson, D. B. (1990), 'ARCH Models as Diffusion Approximations', *Journal of Econometrics* **45**, 7–38.

Ritchken, P. & Trevor, R. (1999), 'Pricing Options under Generalized GARCH and Stochastic Volatility Processes', *Journal of Finance* **59**(1), 377–402.

Scott, L. O. (1987), 'Option Pricing when the Variance Changes Randomly: Theory, Estimation, and an Application', *Journal of Financial and Quantitative Analysis* **22**(4), 419–438.

Stein, E. M. & Stein, J. C. (1991), 'Stock Price Distributions with Stochastic Volatility: An Analytical Approach', *Review of Financial Studies* **4**(4), 727–752.

Stentoft, L. (2004), 'Convergence of the Least Squares Monte Carlo Approach to American Option Valuation', *Management Science* **50**(9), 1193–1203.

Stentoft, L. (2005), 'Pricing American Options when the Underlying Asset follows GARCH Processes', *Journal of Empirical Finance* **12**(4), 576–611.

Stentoft, L. (2008), 'American Option Pricing using GARCH models and the Normal Inverse Gaussian Distribution', *Journal of Financial Econometrics* **6**(4), 540–582.

Tsitsiklis, J. N. & Van Roy, B. (2001), 'Regression Methods for Pricing Complex American-Style Options', *IEEE Transactions on Neural Networks* **12**(4), 694–703.

Wiggins, J. B. (1987), 'Option Values under Stochastic Volatility: Theory and Empirical Estimates', *Journal of Financial Economics* **19**, 351–372.

A Tables

Table 1: Sample statistics for return series

Ticker	GM	IBM	MRK	OEX
Mean	0.0366	0.0250	0.0669	0.0531
Std. Dev.	1.6455	1.4677	1.4469	0.9896
Skewness Statistic	−0.344	−0.940	−0.001	−2.428
	[0.0000]	[0.0000]	[0.9879]	[0.0000]
Ex. Kurt. Statistic	10.761	23.931	3.417	60.987
	[0.0000]	[0.0000]	[0.0000]	[0.0000]
Normality Statistic	24490	121370	2459.8	788375
	[0.0000]	[0.0000]	[0.0000]	[0.0000]

Notes: This table shows sample statistics for the continuously compounded returns, R_t, for the individual stocks and the index considered. The sample period is January 2, 1976, to December 29, 1995, for a total of 5055 observations. For the skewness and excess kurtosis statistics, the brackets below the statistics report the p-values from testing the significance of the difference between the empirical values and the theoretical values from the Normal distribution using a t-test. For the normality statistic the p-value of a t-version of the well known Jarque-Bera test for normality is reported in brackets below the statistics.

Table 2: Estimation results for GM

	Discrete time models							Diffusion limit models			
Model	CV		GARCH		NGARCH			C-GARCH		C-NGARCH	
Loglik	-9690.36		-9360.67		-9337.06			-9360.67		-9337.06	
	Estim.	Std.Err.	Estim.	Std.Err.	Estim.	Std.Err.		Estim.	Std.Err.	Estim.	Std.Err.
λ	0.0174	(0.0141)	0.0322	(0.0140)	0.0134	(0.0140)	λ	0.0322	(0.0140)	0.0134	(0.0140)
ω	2.7077	(0.1360)	0.0305	(0.0207)	0.0191	(0.0102)	β_0	0.0305	(0.0207)	0.0191	(0.0102)
β			0.9330	(0.0303)	0.9392	(0.0196)	β_1	-0.0095	(0.0063)	-0.0057	(0.0035)
α			0.0574	(0.0258)	0.0417	(0.0149)	β_3	0.0812	(0.0365)	0.0589	(0.0211)
γ					-0.5668	(0.1233)	β_2			-0.0472	(0.0162)
	Stat.	P-value	Stat.	P-value	Stat.	P-value		Stat.	P-value	Stat.	P-value
J-B	24490	[0.0000]	1815.7	[0.0000]	1262.6	[0.0000]		1815.7	[0.0000]	1262.6	[0.0000]
Q(20)	41.623	[0.0031]	34.200	[0.0248]	32.802	[0.0355]		34.200	[0.0248]	32.802	[0.0355]
$Q^2(20)$	635.02	[0.0000]	37.661	[0.0043]	36.893	[0.0054]		37.661	[0.0043]	36.893	[0.0054]
ARCH5	116.36	[0.0000]	4.3440	[0.0006]	3.4260	[0.0043]		4.3440	[0.0006]	3.4260	[0.0043]
SIC	3.8342		3.7041		3.6948			3.7041		3.6948	

Notes: This table reports Quasi Maximum Likelihood Estimates (QMLE) for the daily returns assuming a risk-free interest rate of 5.4% corresponding to the value on December 29, 1995. Robust standard errors are reported in parentheses. J-B is the value of the usual Jarque-Bera normality test for the standardized residuals. $Q(20)$ is the Ljung-Box portmanteau test for up to 20'th order serial correlation in the standardized residuals, whereas $Q^2(20)$ is for up to 20'th order serial correlation in the squared standardized residuals. Finally, ARCH5 denotes the ARCH test from Engle (1982). P-values are reported in square brakets. The last row reports the Schwarz Information Criteria.

Table 3: Estimation results for IBM

	Discrete time models							Diffusion limit models			
Model	CV		GARCH		NGARCH			C-GARCH		C-NGARCH	
Loglik	-9112.36		-8798.75		-8769.71			-8798.75		-8769.71	
	Estim.	Std.Err.	Estim.	Std.Err.	Estim.	Std.Err.		Estim.	Std.Err.	Estim.	Std.Err.
λ	0.0098	(0.0140)	0.0316	(0.0171)	0.0100	(0.0140)	λ	0.0316	(0.0171)	0.0100	(0.0140)
ω	2.1542	(0.1543)	0.0246	(0.0143)	0.0270	(0.0131)	β_0	0.0246	(0.0143)	0.0270	(0.0131)
β			0.9391	(0.0266)	0.9265	(0.0271)	β_1	-0.0087	(0.0064)	-0.0104	(0.0061)
α			0.0522	(0.0258)	0.0487	(0.0180)	β_3	0.0738	(0.0365)	0.0689	(0.0254)
γ					-0.5433	(0.1230)	β_2			-0.0529	(0.0262)
	Stat.	P-value	Stat.	P-value	Stat.	P-value		Stat.	P-value	Stat.	P-value
J-B	121373	[0.0000]	13673	[0.0000]	7077.5	[0.0000]		13673	[0.0000]	7077.5	[0.0000]
Q(20)	48.297	[0.0004]	24.823	[0.2083]	24.900	[0.2053]		24.823	[0.2083]	24.900	[0.2053]
$Q^2(20)$	207.82	[0.0000]	9.1608	[0.9559]	10.164	[0.9264]		9.1608	[0.9559]	10.164	[0.9264]
ARCH5	33.512	[0.0000]	0.4413	[0.8199]	0.3919	[0.8547]		0.4413	[0.8199]	0.3919	[0.8547]
SIC	3.6056		3.4818		3.4704			3.4818		3.4704	

Notes: See Table 2.

Table 4: Estimation results for MRK

Model	Discrete time models							Diffusion limit models			
	CV		GARCH		NGARCH			C-GARCH		C-NGARCH	
Loglik	-9040.23		-8858.53		-8849.72			-8858.53		-8849.72	
	Estim.	Std.Err.	Estim.	Std.Err.	Estim.	Std.Err.		Estim.	Std.Err.	Estim.	Std.Err.
λ	0.0386	(0.0141)	0.0512	(0.0139)	0.0407	(0.0141)	λ	0.0512	(0.0139)	0.0407	(0.0141)
ω	2.0936	(0.0685)	0.0692	(0.0281)	0.0628	(0.0198)	β_0	0.0692	(0.0281)	0.0628	(0.0198)
β			0.9072	(0.0276)	0.9100	(0.0205)	β_1	-0.0328	(0.0131)	-0.0300	(0.0094)
α			0.0600	(0.0169)	0.0535	(0.0128)	β_3	0.0848	(0.0239)	0.0756	(0.0181)
γ					-0.3505	(0.1211)	β_2			-0.0375	(0.0139)
	Stat.	P-value	Stat.	P-value	Stat.	P-value		Stat.	P-value	Stat.	P-value
J-B	2459.8	[0.0000]	523.78	[0.0000]	431.80	[0.0000]		523.78	[0.0000]	431.80	[0.0000]
Q(20)	38.619	[0.0074]	38.847	[0.0070]	39.009	[0.0067]		38.847	[0.0070]	39.009	[0.0067]
$Q^2(20)$	507.39	[0.0000]	20.024	[0.3315]	22.846	[0.1966]		20.024	[0.3315]	22.846	[0.1966]
ARCH5	35.317	[0.0000]	1.4345	[0.2083]	1.8074	[0.1078]		1.4345	[0.2083]	1.8074	[0.1078]
SIC	3.5770		3.5054		3.5020			3.5054		3.5020	

Notes: See Table 2.

Table 5: Estimation results for OEX

Model	Discrete time models							Diffusion limit models			
	CV		GARCH		NGARCH			C-GARCH		C-NGARCH	
Loglik	-7119.99		-6519.31		-6503.83			-6519.31		-6503.83	
	Estim.	Std.Err.	Estim.	Std.Err.	Estim.	Std.Err.		Estim.	Std.Err.	Estim.	Std.Err.
λ	0.0370	(0.0146)	0.0529	(0.0150)	0.0380	(0.0142)	λ	0.0529	(0.0150)	0.0380	(0.0142)
ω	0.9793	(0.1093)	0.0112	(0.0065)	0.0124	(0.0070)	β_0	0.0112	(0.0065)	0.0124	(0.0070)
β			0.9364	(0.0274)	0.9286	(0.0299)	β_1	-0.0113	(0.0060)	-0.0129	(0.0068)
α			0.0523	(0.0246)	0.0488	(0.0192)	β_3	0.0740	(0.0347)	0.0690	(0.0272)
γ					-0.4481	(0.1188)	β_2			-0.0437	(0.0252)
	Stat.	P-value	Stat.	P-value	Stat.	P-value		Stat.	P-value	Stat.	P-value
J-B	788375	[0.0000]	9200.3	[0.0000]	7088.0	[0.0000]		9200.3	[0.0000]	7088.0	[0.0000]
Q(20)	47.095	[0.0006]	32.408	[0.0391]	29.014	[0.0875]		32.408	[0.0391]	29.014	[0.0875]
$Q^2(20)$	435.65	[0.0000]	16.775	[0.5386]	12.695	[0.8093]		16.775	[0.5386]	12.695	[0.8093]
ARCH5	66.323	[0.0000]	1.5535	[0.1697]	0.9175	[0.4683]		1.5535	[0.1697]	0.9175	[0.4683]
SIC	2.8173		2.5799		2.5739			2.5799		2.5739	

Notes: See Table 2.

Table 6: Parameter values used for the GARCH models

Mean specification	
Equation	r (annualized)
$S_t = S_{t-1} \times \exp\left\{r - \frac{1}{2}h_t + \sqrt{h_t}\varepsilon_t^*\right\}$	0.06

Volatility specifications					
Model	Equation	ω	β	α	γ
CV	$h_t = \omega$	$2.48 * 10^{-4}$			
GARCH	$h_t = \omega + \beta h_{t-1} + \alpha h_{t-1}\left(\tilde{\varepsilon}_{t-1}^*\right)^2$	$4.96 * 10^{-6}$	0.92	0.060	
NGARCH	$h_t = \omega + \beta h_{t-1} + \alpha h_{t-1}\left(\tilde{\varepsilon}_{t-1}^* + \gamma\right)^2$	$4.96 * 10^{-6}$	0.92	0.048	-0.5

Notes: This table reports the parameter values used in the Monte Carlo study of the discrete time GARCH models. The sample paths are generated with $\varepsilon_t^* \sim N(0,1)$ and $\tilde{\varepsilon}_t^* = \varepsilon_t^* - \lambda$, with $\lambda = 0.05$. In all specifications the initial level of the volatility is set equal to the unconditional level.

Table 7: Parameter values used for the GARCH diffusion limit models

Mean specification	
Equation	r (annualized)
$d\ln S_t = \left(r - \frac{1}{2}h_t\right)dt + \sqrt{h_t}dW_{1t}^*$	0.06

Volatility specifications					
Model	Equation	β_0	β_1	β_2	β_3
C-GARCH	$dh_t = \beta_0 dt + \beta_1 h_t dt + \beta_3 h_t dW_{2t}^*$	$4.96 * 10^{-6}$	-0.032		0.0679
C-NGARCH	$dh_t = \beta_0 dt + \beta_1 h_t dt + \beta_2 h_t d\tilde{W}_{1t}^* + \beta_3 h_t dW_{2t}^*$	$4.96 * 10^{-6}$	-0.020	-0.048	0.0679

Notes: This table reports the parameter values used in the Monte Carlo study of the GARCH diffusion limit models. The sample paths are generated with W_{1t}^* and W_{2t}^* as two independent Wiener proceses and $d\tilde{W}_{1t}^* = dW_{1t}^* - \lambda dt$, with $\lambda = 0.05$. In all specifications the initial level of the volatility is set equal to the unconditional level.

Table 8: American put price estimates in models with different volatility processes

		CV	GARCH		NGARCH		C-GARCH		C-NGARCH	
T	K	Price	Price	RBIAS	Price	RBIAS	Price	RBIAS	Price	RBIAS
7	80	0.000 (0.0000)	0.000 (0.0000)	−97.82	0.000 (0.0000)	−98.82	0.000 (0.0000)	−46.63	0.000 (0.0000)	−94.98
7	90	0.007 (0.0004)	0.011 (0.0005)	−42.37	0.014 (0.0006)	−51.89	0.007 (0.0004)	−8.85	0.013 (0.0005)	−49.63
7	100	1.587 (0.0067)	1.562 (0.0071)	1.65	1.536 (0.0073)	3.35	1.586 (0.0068)	0.11	1.592 (0.0071)	−0.31
7	110	9.976 (0.0057)	9.978 (0.0059)	−0.02	9.975 (0.0051)	0.01	9.977 (0.0055)	−0.01	9.976 (0.0052)	0.00
7	120	19.972 (0.0053)	19.972 (0.0050)	0.00	19.972 (0.0049)	0.00	19.973 (0.0051)	0.00	19.972 (0.0050)	0.00
21	80	0.001 (0.0002)	0.006 (0.0005)	−74.05	0.011 (0.0008)	−86.25	0.002 (0.0003)	−38.76	0.009 (0.0006)	−82.56
21	90	0.189 (0.0031)	0.210 (0.0032)	−9.96	0.252 (0.0037)	−24.93	0.195 (0.0030)	−2.89	0.265 (0.0038)	−28.75
21	100	2.662 (0.0108)	2.610 (0.0115)	2.03	2.591 (0.0118)	2.76	2.652 (0.0119)	0.38	2.683 (0.0125)	−0.75
21	110	10.085 (0.0117)	10.078 (0.0112)	0.07	10.026 (0.0095)	0.59	10.092 (0.0118)	−0.07	10.043 (0.0106)	0.42
21	120	19.972 (0.0053)	19.972 (0.0051)	0.00	19.972 (0.0049)	0.00	19.973 (0.0051)	0.00	19.973 (0.0050)	0.00
63	80	0.123 (0.0026)	0.165 (0.0034)	−25.22	0.263 (0.0047)	−53.22	0.140 (0.0030)	−11.84	0.263 (0.0045)	−53.11
63	90	1.072 (0.0085)	1.082 (0.0095)	−0.93	1.260 (0.0113)	−14.92	1.072 (0.0080)	0.00	1.310 (0.0097)	−18.18
63	100	4.364 (0.0176)	4.268 (0.0202)	2.23	4.326 (0.0211)	0.86	4.328 (0.0142)	0.81	4.449 (0.0152)	−1.93
63	110	10.865 (0.0218)	10.786 (0.0204)	0.73	10.652 (0.0198)	2.00	10.861 (0.0197)	0.04	10.757 (0.0204)	1.00
63	120	19.980 (0.0086)	19.984 (0.0078)	−0.02	19.974 (0.0052)	0.03	19.987 (0.0091)	−0.04	19.975 (0.0055)	0.03
126	80	0.547 (0.0066)	0.602 (0.0078)	−9.13	0.863 (0.0104)	−36.56	0.567 (0.0070)	−3.41	0.872 (0.0094)	−37.26
126	90	2.183 (0.0142)	2.171 (0.0147)	0.56	2.498 (0.0172)	−12.59	2.167 (0.0131)	0.72	2.563 (0.0157)	−14.81
126	100	5.843 (0.0222)	5.730 (0.0236)	1.96	5.922 (0.0253)	−1.34	5.790 (0.0205)	0.91	6.049 (0.0224)	−3.40
126	110	11.901 (0.0294)	11.767 (0.0275)	1.14	11.715 (0.0296)	1.59	11.864 (0.0267)	0.31	11.859 (0.0283)	0.36
126	120	20.174 (0.0229)	20.130 (0.0192)	0.22	20.042 (0.0162)	0.66	20.184 (0.0225)	−0.05	20.090 (0.0200)	0.42

Notes: This table shows American put prices for the set of artificial options with T denoting the time to maturity in days and K denoting the strike price. The initial stock price is set to 100. The parameter values for the different GARCH processes are the ones specified in the text and in Table 6 for the discrete time models and in Table 7 for the diffusion models. For the diffusion limits an Euler discretization with 1024 intraday steps is used. In the cross-sectional regressions powers of and cross products between the stock level and the level of the volatility of total order less than or equal to three were used. Exercise is considered once every trading day. Prices reported are averages of 100 calculated prices using 100,000 paths and different seeds in the random number generator. In parentheses standard errors of these price estimates are reported. The column headed RBIAS reports the difference between the CV model and the time varying volatility model relative to the latter in percentage terms. Thus, it indicates the relative mispricing by the CV model which would be observed if the true model has the corresponding GARCH specification.

Table 9: European put price estimates in models with different volatility processes

		CV	GARCH		NGARCH		C-GARCH		C-NGARCH	
T	K	Price	Price	RBIAS	Price	RBIAS	Price	RBIAS	Price	RBIAS
7	80	0.000 (0.0000)	0.000 (0.0000)	−97.82	0.000 (0.0000)	−98.82	0.000 (0.0000)	−46.63	0.000 (0.0000)	−94.98
7	90	0.006 (0.0004)	0.011 (0.0005)	−41.72	0.013 (0.0006)	−51.52	0.007 (0.0004)	−8.05	0.013 (0.0006)	−49.41
7	100	1.579 (0.0076)	1.553 (0.0076)	1.67	1.528 (0.0076)	3.36	1.576 (0.0076)	0.20	1.583 (0.0079)	−0.24
7	110	9.837 (0.0133)	9.842 (0.0149)	−0.05	9.831 (0.0147)	0.06	9.838 (0.0140)	−0.01	9.830 (0.0141)	0.07
7	120	19.802 (0.0135)	19.801 (0.0151)	0.01	19.801 (0.0147)	0.01	19.802 (0.0141)	0.00	19.802 (0.0141)	0.00
21	80	0.001 (0.0002)	0.005 (0.0006)	−73.80	0.010 (0.0009)	−86.40	0.002 (0.0003)	−36.74	0.008 (0.0006)	−82.75
21	90	0.187 (0.0032)	0.207 (0.0035)	−9.71	0.249 (0.0040)	−24.75	0.192 (0.0030)	−2.51	0.262 (0.0038)	−28.52
21	100	2.628 (0.0119)	2.580 (0.0125)	1.89	2.563 (0.0129)	2.55	2.615 (0.0124)	0.50	2.648 (0.0132)	−0.72
21	110	9.832 (0.0193)	9.832 (0.0204)	0.00	9.747 (0.0208)	0.88	9.835 (0.0210)	−0.03	9.763 (0.0218)	0.71
21	120	19.420 (0.0214)	19.429 (0.0225)	−0.05	19.411 (0.0225)	0.05	19.423 (0.0236)	−0.01	19.408 (0.0239)	0.06
63	80	0.121 (0.0028)	0.160 (0.0038)	−24.77	0.257 (0.0052)	−53.10	0.136 (0.0032)	−11.20	0.256 (0.0050)	−52.99
63	90	1.050 (0.0095)	1.062 (0.0100)	−1.14	1.239 (0.0116)	−15.25	1.048 (0.0093)	0.19	1.284 (0.0115)	−18.19
63	100	4.237 (0.0216)	4.159 (0.0214)	1.89	4.225 (0.0228)	0.28	4.198 (0.0184)	0.93	4.326 (0.0204)	−2.05
63	110	10.405 (0.0332)	10.347 (0.0312)	0.56	10.212 (0.0328)	1.90	10.385 (0.0276)	0.20	10.284 (0.0295)	1.18
63	120	18.779 (0.0393)	18.791 (0.0361)	−0.06	18.632 (0.0374)	0.79	18.793 (0.0331)	−0.07	18.643 (0.0348)	0.73
126	80	0.531 (0.0065)	0.584 (0.0089)	−9.04	0.839 (0.0121)	−36.70	0.548 (0.0076)	−3.00	0.845 (0.0107)	−37.17
126	90	2.104 (0.0157)	2.098 (0.0172)	0.31	2.421 (0.0202)	−13.09	2.085 (0.0153)	0.93	2.471 (0.0184)	−14.85
126	100	5.560 (0.0274)	5.472 (0.0276)	1.61	5.675 (0.0302)	−2.03	5.504 (0.0235)	1.01	5.767 (0.0267)	−3.60
126	110	11.143 (0.0397)	11.043 (0.0382)	0.91	11.015 (0.0408)	1.16	11.093 (0.0316)	0.45	11.106 (0.0345)	0.33
126	120	18.501 (0.0478)	18.452 (0.0462)	0.26	18.266 (0.0490)	1.29	18.485 (0.0383)	0.08	18.318 (0.0408)	1.00

Notes: This table shows European put prices for the set of artificial options. See also the notes to Table 8.

Table 10: Early exercise value estimates in models with different volatility processes

		CV	GARCH		NGARCH		C-GARCH		C-NGARCH	
T	K	EE val	EE val	RBIAS	EE val	RBIAS	EE val	RBIAS	EE val	RBIAS
7	80	0.000 (0.0000)	0.000 (0.0000)	0.00	0.000 (0.0000)	0.00	0.000 (0.0000)	0.00	0.000 (0.0000)	0.00
7	90	0.000 (0.0001)	0.000 (0.0001)	−67.64	0.000 (0.0001)	−68.69	0.000 (0.0001)	−43.44	0.000 (0.0001)	−61.47
7	100	0.008 (0.0030)	0.008 (0.0030)	−3.33	0.008 (0.0031)	2.27	0.009 (0.0035)	−15.70	0.009 (0.0036)	−13.22
7	110	0.139 (0.0118)	0.136 (0.0132)	1.93	0.144 (0.0136)	−3.83	0.139 (0.0125)	0.14	0.146 (0.0126)	−4.67
7	120	0.170) (0.0122	0.171 (0.0139)	−0.74	0.171 (0.0136)	−0.82	0.171 (0.0128)	−0.28	0.171 (0.0129)	−0.27
21	80	0.000 (0.0001)	0.000 (0.0002)	−76.74	0.001 (0.0003)	−84.22	0.000 (0.0001)	−55.82	0.001 (0.0002)	−79.76
21	90	0.002 (0.0009)	0.003 (0.0012)	−29.46	0.003 (0.0015)	−39.39	0.003 (0.0011)	−29.81	0.003 (0.0013)	−45.91
21	100	0.034 (0.0054)	0.030 (0.0057)	13.70	0.028 (0.0067)	21.70	0.037 (0.0060)	−7.88	0.035 (0.0065)	−2.80
21	110	0.253 (0.0151)	0.246 (0.0161)	2.90	0.280 (0.0183)	−9.74	0.257 (0.0180)	−1.59	0.280 (0.0205)	−9.74
21	120	0.552 (0.0202)	0.543 (0.0222)	1.66	0.561 (0.0222)	−1.61	0.550 (0.0226)	0.36	0.564 (0.0229)	−2.20
63	80	0.003 (0.0012)	0.005 (0.0014)	−41.14	0.006 (0.0022)	−57.94	0.004 (0.0012)	−33.39	0.006 (0.0022)	−57.79
63	90	0.022 (0.0043)	0.020 (0.0049)	10.23	0.021 (0.0062)	4.88	0.024 (0.0041)	−8.22	0.026 (0.0051)	−17.77
63	100	0.126 (0.0123)	0.110 (0.0124)	15.14	0.101 (0.0136)	24.84	0.130 (0.0110)	−3.08	0.124 (0.0124)	2.27
63	110	0.460 (0.0231)	0.439 (0.0249)	4.81	0.440 (0.0274)	4.41	0.476 (0.0209)	−3.39	0.473 (0.0226)	−2.85
63	120	1.200 (0.0378)	1.192 (0.0371)	0.67	1.342 (0.0375)	−10.53	1.194 (0.0328)	0.53	1.331 (0.0344)	−9.83
126	80	0.016 (0.0036)	0.018 (0.0040)	−12.05	0.023 (0.0052)	−31.53	0.019 (0.0041)	−15.18	0.027 (0.0054)	−40.11
126	90	0.079 (0.0080)	0.073 (0.0093)	7.92	0.076 (0.0121)	3.24	0.082 (0.0087)	−4.48	0.091 (0.0109)	−13.70
126	100	0.283 (0.0160)	0.259 (0.0162)	9.39	0.247 (0.0184)	14.47	0.286 (0.0181)	−0.94	0.282 (0.0199)	0.54
126	110	0.758 (0.0270)	0.724 (0.0275)	4.72	0.700 (0.0302)	8.27	0.772 (0.0266)	−1.74	0.752 (0.0299)	0.79
126	120	1.674 (0.0424)	1.678 (0.0446)	−0.25	1.776 (0.0492)	−5.77	1.699 (0.0341)	−1.52	1.772 (0.0374)	−5.54

Notes: This table shows early exercise values for the set of artificial put options. See also the notes to Table 8.

Table 11: Overall performance in terms of dollar errors

Panel A: GM

	Discrete time models				Diffusion limit models			
Model	Put	(629)	Call	(1206)	Put	(629)	Call	(1206)
	BIAS	RMSE	BIAS	RMSE	BIAS	RMSE	BIAS	RMSE
CV	-0.205	0.312	-0.219	0.315	-0.205	0.313	-0.223	0.317
GARCH	-0.019	0.222	-0.046	0.226	-0.009	0.231	-0.035	0.228
NGARCH	-0.001	0.217	-0.057	0.223	0.005	0.226	-0.051	0.228

Panel B: IBM

	Discrete time models				Diffusion limit models			
Model	Put	(1827)	Call	(2918)	Put	(1827)	Call	(2918)
	BIAS	RMSE	BIAS	RMSE	BIAS	RMSE	BIAS	RMSE
CV	-0.323	0.528	-0.384	0.604	-0.325	0.529	-0.391	0.608
GARCH	-0.074	0.514	-0.117	0.540	-0.094	0.466	-0.127	0.488
NGARCH	-0.065	0.524	-0.203	0.557	-0.079	0.474	-0.218	0.510

Panel C: MRK

	Discrete time models				Diffusion limit models			
Model	Put	(553)	Call	(1291)	Put	(553)	Call	(1291)
	BIAS	RMSE	BIAS	RMSE	BIAS	RMSE	BIAS	RMSE
CV	-0.194	0.344	-0.188	0.385	-0.194	0.344	-0.194	0.388
GARCH	-0.161	0.299	-0.164	0.339	-0.163	0.304	-0.159	0.342
NGARCH	-0.137	0.271	-0.157	0.320	-0.137	0.274	-0.151	0.320

Panel D: OEX

	Discrete time models				Diffusion limit models			
Model	Put	(4804)	Call	(3487)	Put	(4804)	Call	(3487)
	BIAS	RMSE	BIAS	RMSE	BIAS	RMSE	BIAS	RMSE
CV	0.744	1.591	1.750	2.274	0.739	1.588	1.710	2.244
GARCH	-0.411	0.916	0.405	0.975	-0.414	0.942	0.428	1.001
NGARCH	-0.314	0.855	0.353	0.929	-0.306	0.867	0.367	0.946

Notes: This table shows the performance of the discrete time GARCH models and their diffusion limits in terms of dollar errors. We report results for the two metrics described in the text. Thus, denoting the k'th price estimate by $\bar{P}_k$ and the k'th observed price by P_k these are the bias, $BIAS \equiv K^{-1}\sum_{k=1}^{K}\left(\bar{P}_k - P_k\right)$, and the root mean squared error, $RMSE \equiv \sqrt{K^{-1}\sum_{k=1}^{K}\left(\bar{P}_k - P_k\right)^2}$.

Table 12: Overall performance in terms of ISD errors

Panel A: GM

Model	Discrete time models Put	(629)	Call	(1206)	Diffusion limit models Put	(629)	Call	(1206)
	BIAS	RMSE	BIAS	RMSE	BIAS	RMSE	BIAS	RMSE
CV	-3.973%	8.086%	-5.913%	12.591%	-3.973%	8.086%	-5.913%	12.591%
GARCH	0.157%	6.586%	-1.960%	10.492%	0.312%	6.728%	-1.788%	10.539%
NGARCH	0.537%	6.462%	-2.113%	10.663%	0.634%	6.552%	-1.865%	10.288%

Panel B: IBM

Model	Discrete time models Put	(1827)	Call	(2918)	Diffusion limit models Put	(1827)	Call	(2918)
	BIAS	RMSE	BIAS	RMSE	BIAS	RMSE	BIAS	RMSE
CV	-4.245%	8.033%	-5.498%	9.710%	-4.245%	8.033%	-5.498%	9.710%
GARCH	-0.927%	6.495%	-2.186%	7.343%	-1.166%	6.383%	-2.294%	7.334%
NGARCH	-0.810%	6.404%	-2.892%	7.371%	-0.945%	6.225%	-3.036%	7.403%

Panel C: MRK

Model	Discrete time models Put	(553)	Call	(1291)	Diffusion limit models Put	(553)	Call	(1291)
	BIAS	RMSE	BIAS	RMSE	BIAS	RMSE	BIAS	RMSE
CV	-3.837%	7.860%	-4.111%	9.980%	-3.837%	7.860%	-4.111%	9.980%
GARCH	-2.844%	6.992%	-3.281%	8.858%	-3.078%	7.797%	-3.293%	8.960%
NGARCH	-2.479%	6.734%	-3.133%	8.564%	-2.681%	7.519%	-3.154%	8.753%

Panel D: OEX

Model	Discrete time models Put	(4804)	Call	(3487)	Diffusion limit models Put	(4804)	Call	(3487)
	BIAS	RMSE	BIAS	RMSE	BIAS	RMSE	BIAS	RMSE
CV	0.522%	5.111%	3.601%	5.213%	0.522%	5.111%	3.601%	5.213%
GARCH	-2.335%	4.762%	0.488%	3.377%	-2.554%	5.376%	0.541%	3.455%
NGARCH	-1.968%	4.506%	0.326%	3.343%	-2.044%	4.768%	0.350%	3.352%

Notes: This table shows the performance of the discrete time GARCH models and their diffusion limits in terms of ISD errors. We report results for the two metrics described in the text and in the notes to Table 11.

Table 13: Performance in terms of ISD errors for ITM and D-ITM put options with one intraday early exercise

Panel A: GM

Model	Discrete time models				Diffusion limit models			
	ITM	(103)	D-ITM	(103)	ITM	(103)	D-ITM	(103)
	BIAS	RMSE	BIAS	RMSE	BIAS	RMSE	BIAS	RMSE
CV	-3.27%	4.94%	-1.31%	13.50%	-3.27%	4.94%	-1.31%	13.50%
GARCH	-0.51%	3.22%	4.16%	13.07%	-0.12%	3.32%	4.31%	13.12%
NGARCH	-0.54%	3.90%	3.19%	12.97%	-0.41%	4.01%	2.97%	12.94%

Panel B: IBM

Model	Discrete time models				Diffusion limit models			
	ITM	(290)	D-ITM	(168)	ITM	(290)	D-ITM	(168)
	BIAS	RMSE	BIAS	RMSE	BIAS	RMSE	BIAS	RMSE
CV	-3.40%	6.12%	-3.55%	14.91%	-3.40%	6.12%	-3.55%	14.91%
GARCH	-0.32%	5.41%	3.16%	14.05%	-0.11%	5.19%	1.55%	13.45%
NGARCH	-1.01%	5.57%	2.01%	14.19%	-0.88%	5.34%	0.54%	13.71%

Panel C: MRK

Model	Discrete time models				Diffusion limit models			
	ITM	(61)	D-ITM	(70)	ITM	(61)	D-ITM	(70)
	BIAS	RMSE	BIAS	RMSE	BIAS	RMSE	BIAS	RMSE
CV	-2.41%	4.54%	-7.27%	16.42%	-2.41%	4.54%	-7.27%	16.42%
GARCH	-1.22%	3.66%	-5.54%	16.01%	-0.84%	3.61%	-6.80%	18.42%
NGARCH	-1.19%	3.56%	-5.77%	16.02%	-0.91%	3.41%	-7.06%	18.42%

Panel D: ALL

Model	Discrete time models				Diffusion limit models			
	ITM	(454)	D-ITM	(341)	ITM	(454)	D-ITM	(341)
	BIAS	RMSE	BIAS	RMSE	BIAS	RMSE	BIAS	RMSE
CV	-3.24%	5.68%	-3.64%	14.83%	-3.24%	5.68%	-3.64%	14.83%
GARCH	-0.49%	4.78%	1.68%	14.20%	-0.21%	4.63%	0.67%	14.52%
NGARCH	-0.93%	4.99%	0.77%	14.24%	-0.78%	4.84%	-0.29%	14.59%

Notes: This table shows the performance of the discrete time GARCH models and their diffusion limits with one intraday early exercise in terms of ISD errors. We report results for two metrics described in the text and in the notes to Table 11 for the subsample of in the money and deep in the money put options on individual stocks.

Table 14: Performance in terms of ISD errors for ITM and D-ITM options with multiple intraday early exercises

Panel A: GM

	One intraday early exercise				Multiple intraday early exercises			
Model	ITM	(103)	D-ITM	(103)	ITM	(103)	D-ITM	(103)
	BIAS	RMSE	BIAS	RMSE	BIAS	RMSE	BIAS	RMSE
CV	-3.27%	4.94%	-1.31%	13.50%	-3.26%	4.93%	-1.03%	13.58%
GARCH	-0.12%	3.32%	4.31%	13.12%	-0.15%	3.32%	4.39%	13.29%
NGARCH	-0.41%	4.01%	2.97%	12.94%	-0.44%	4.00%	3.13%	13.11%

Panel B: IBM

	One intraday early exercise				Multiple intraday early exercises			
Model	ITM	(290)	D-ITM	(168)	ITM	(290)	D-ITM	(168)
	BIAS	RMSE	BIAS	RMSE	BIAS	RMSE	BIAS	RMSE
CV	-3.40%	6.12%	-3.55%	14.91%	-3.36%	6.10%	-3.43%	14.65%
GARCH	-0.11%	5.19%	1.55%	13.45%	-0.13%	5.17%	1.71%	13.48%
NGARCH	-0.88%	5.34%	0.54%	13.71%	-0.91%	5.32%	0.55%	13.86%

Panel C: MRK

	One intraday early exercise				Multiple intraday early exercises			
Model	ITM	(61)	D-ITM	(70)	ITM	(61)	D-ITM	(70)
	BIAS	RMSE	BIAS	RMSE	BIAS	RMSE	BIAS	RMSE
CV	-2.41%	4.54%	-7.27%	16.42%	-2.38%	4.54%	-6.61%	15.40%
GARCH	-0.84%	3.61%	-6.80%	18.42%	-0.86%	3.62%	-6.51%	18.24%
NGARCH	-0.91%	3.41%	-7.06%	18.42%	-0.93%	3.42%	-6.06%	15.54%

Panel D: ALL

	One intraday early exercise				Multiple intraday early exercises			
Model	ITM	(454)	D-ITM	(341)	ITM	(454)	D-ITM	(341)
	BIAS	RMSE	BIAS	RMSE	BIAS	RMSE	BIAS	RMSE
CV	-3.24%	5.68%	-3.64%	14.83%	-3.20%	5.66%	-3.36%	14.50%
GARCH	-0.21%	4.63%	0.67%	14.52%	-0.23%	4.62%	0.83%	14.53%
NGARCH	-0.78%	4.84%	-0.29%	14.59%	-0.81%	4.83%	-0.03%	14.01%

Notes: This table shows the performance of the diffusion limits of the GARCH models with one and with multiple intraday early exercises in terms of ISD errors. We report results for the two metrics described in the text and in the notes to Table 11 for the subsample of in the money and deep in the money put options on individual stocks.

B Figures

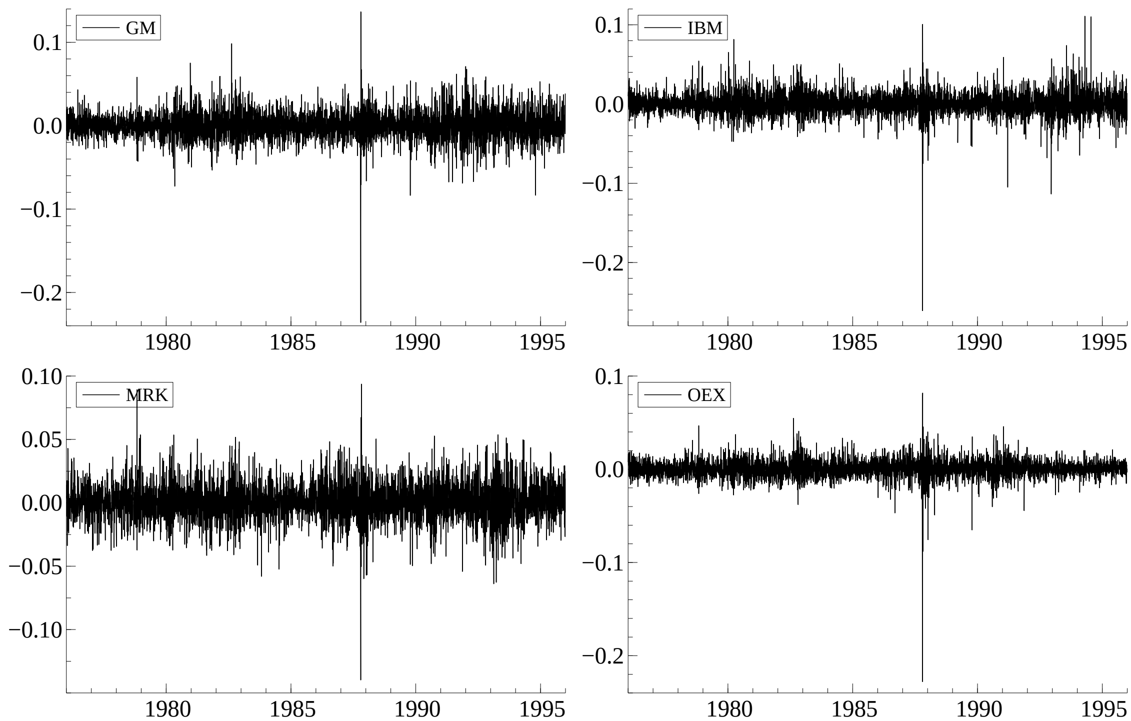

Figure 1: This figure shows time series plots of the annualized continuously compounded returns for the four assets considered. The sample period is January 2, 1976, to December 29, 1995, for a total of 5,055 observations.

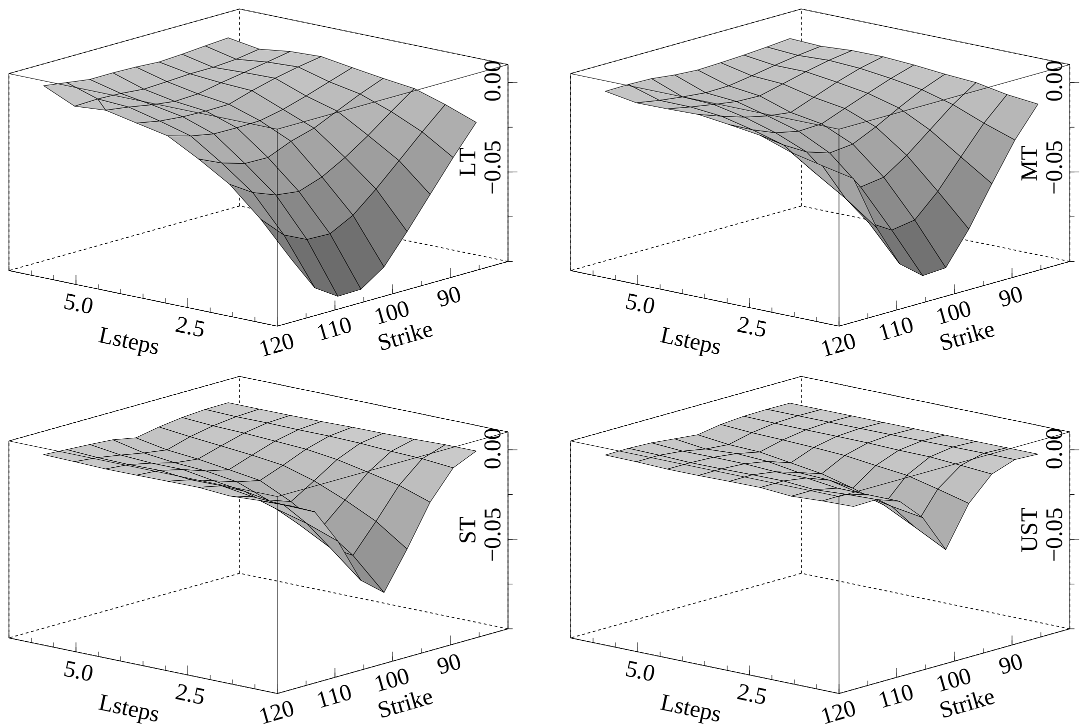

Figure 2: This figure shows the pricing errors for the discrete time models as a function of the strike price and the logarithm of the number of intraday steps for the NGARCH model for the American options with exercise considered at the end of each trading day. The pricing error is the difference between the discrete time model price and the continuous time benchmark price. Each plot is for one maturity and shows the convergence pattern for all strike prices. From left to right and top to bottom the plots are for long term, or LT, options with $T = 126$ days to maturity, middle term, or MT, options with $T = 63$ days to maturity, short term, or ST, options with $T = 21$ days to maturity, and ultra short term, or UST, options with $T = 7$ days to maturity.

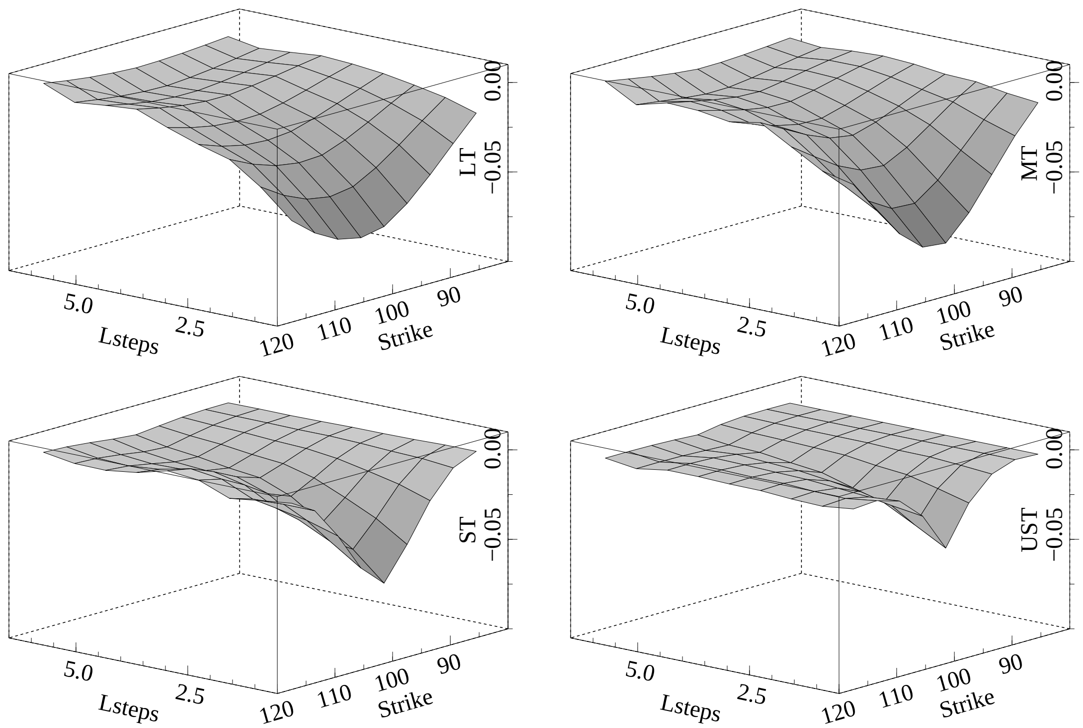

Figure 3: This figure shows the pricing errors for the discrete time models as a function of the strike price and the logarithm of the number of intraday steps for the NGARCH model for the European options. The pricing error is the difference between the discrete time model price and the continuous time benchmark price. Each plot is for one maturity and shows the convergence pattern for all strike prices. From left to right and top to bottom the plots are for long term, or LT, options with $T = 126$ days to maturity, middle term, or MT, options with $T = 63$ days to maturity, short term, or ST, options with $T = 21$ days to maturity, and ultra short term, or UST, options with $T = 7$ days to maturity.

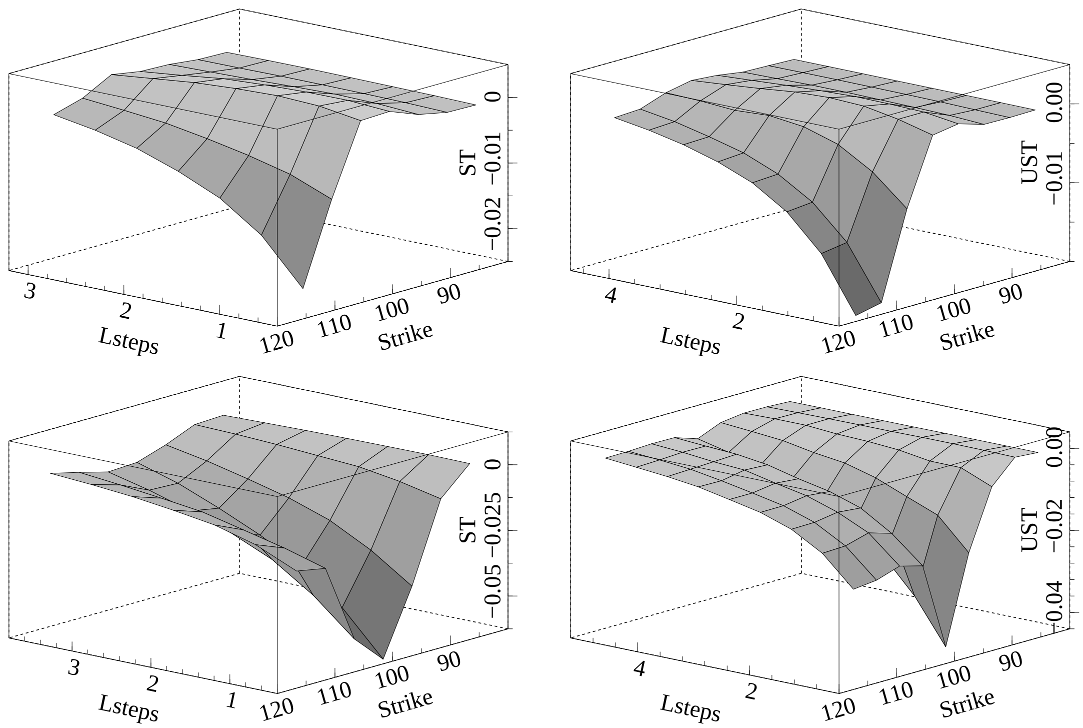

Figure 4: This figure shows the pricing errors for the Bermudan options as a function of the strike price and the logarithm of the number of intraday steps. The pricing error is the difference between the Bermudan option and the American option with maximum number of intraday exercises. In the top two panels only the number of intraday early exercises is increase while the discretization is kept constant. In the bottom two panels the number of intraday discretization is increased along with the number of early exercise opportunities. Left plots are for short term, or ST, options with $T = 21$ days to maturity, and right plots are for ultra short term, or UST, options with $T = 7$ days to maturity.

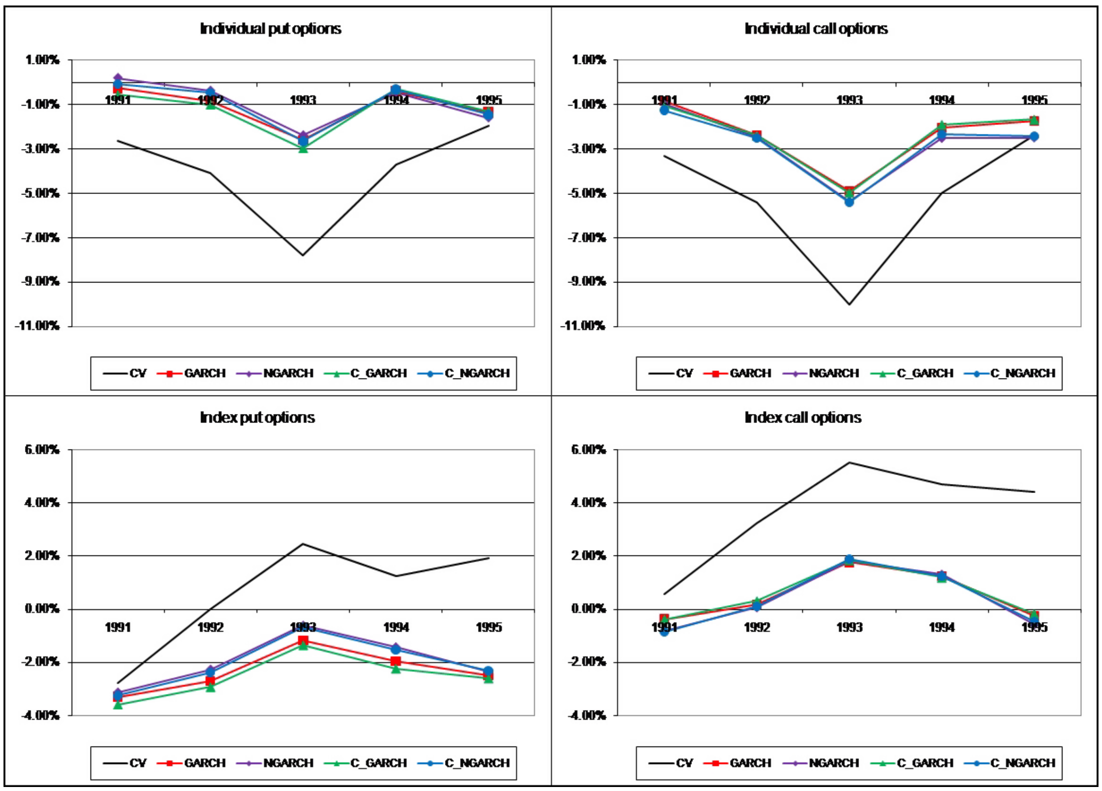

Figure 5: This figure shows the pricing errors through time. The pricing error used is the ISD BIAS as defined in the text. The top row is for individual stock options and the bottom row is for index options. Left hand plots are for put options and right hand plots are for call options.

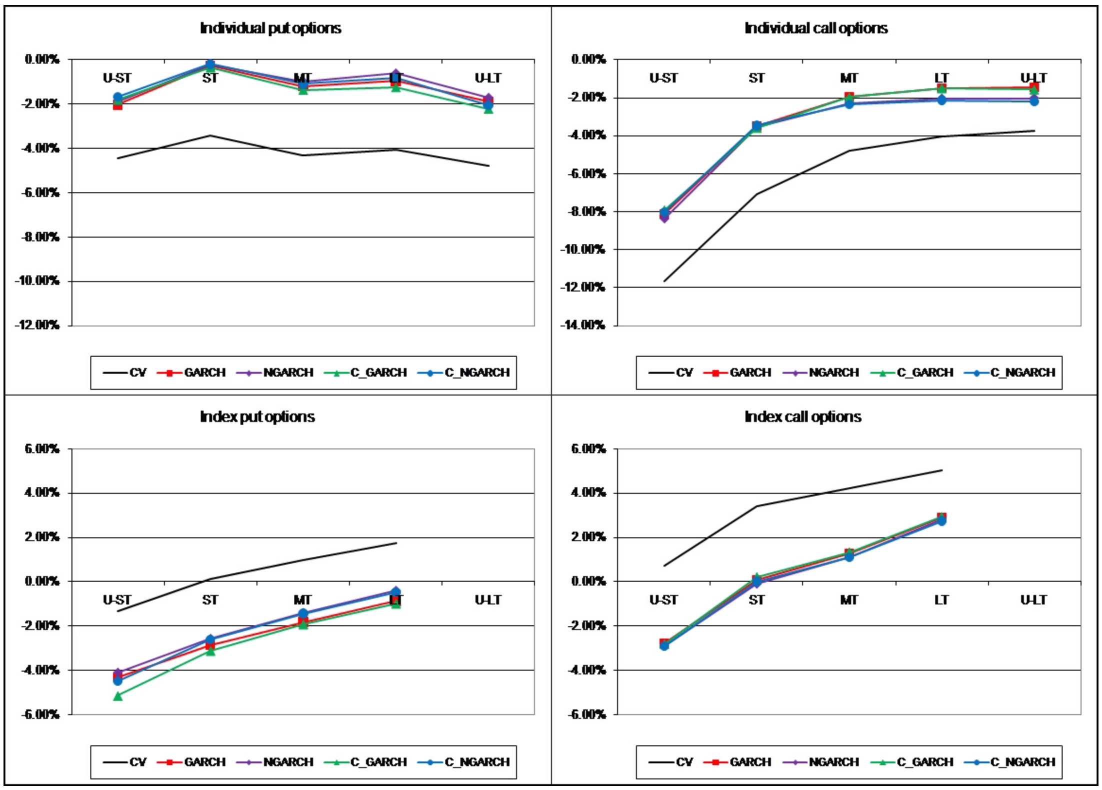

Figure 6: This figure shows the pricing errors across maturity categories. The pricing error used is the ISD BIAS as defined in the text. The top row is for individual stock options and the bottom row is for index options. Left hand plots are for put options and right hand plots are for call options. Note that there are no ultra long term index options available in our sample.

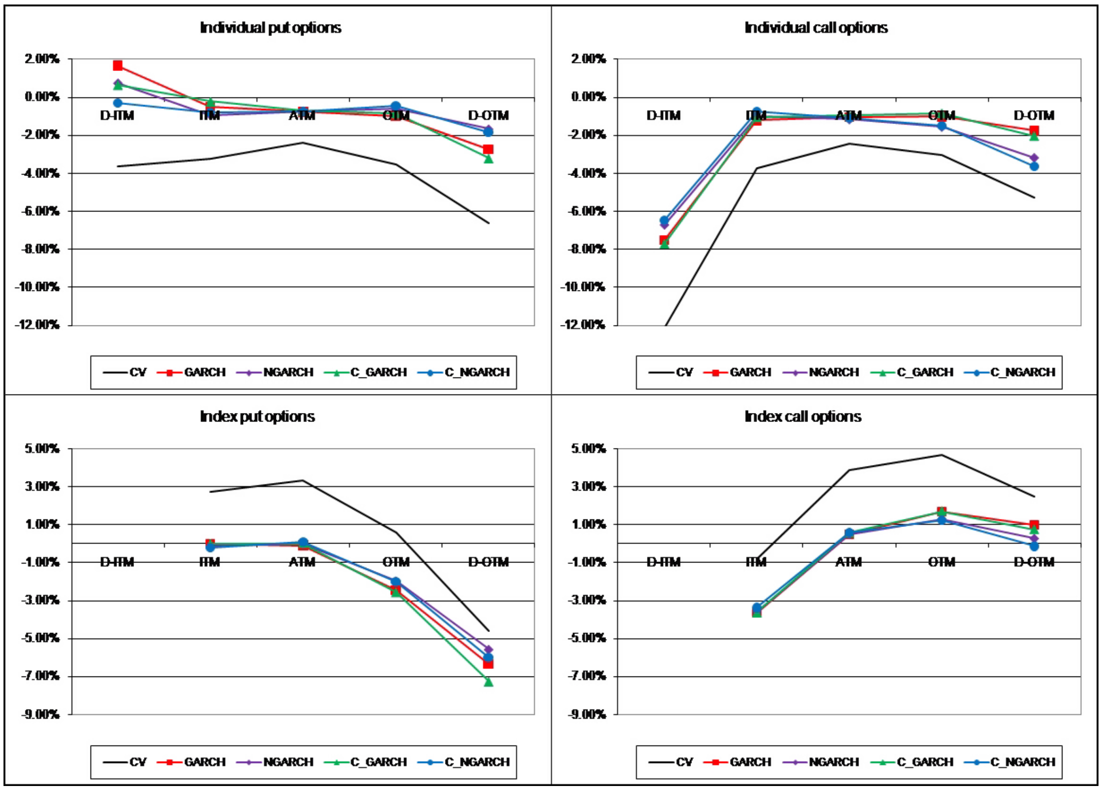

Figure 7: This figure shows the pricing errors across moneyness categories. The pricing error used is the ISD BIAS as defined in the text. The top row is for individual stock options and the bottom row is for index options. Left hand plots are for put options and right hand plots are for call options. Note that there are very few deep in the money index options in our sample and these have been grouped together with the in the money options in the figure.

Research Papers 2007

2011-21: Bent Jesper Christensen, Olaf Posch and Michel van der Wel: Estimating Dynamic Equilibrium Models using Macro and Financial Data

2011-22: Antonis Papapantoleon, John Schoenmakers and David Skovmand: Efficient and accurate log-Lévi approximations to Lévi driven LIBOR models

2011-23: Torben G. Andersen, Dobrislav Dobrev and Ernst Schaumburg: A Functional Filtering and Neighborhood Truncation Approach to Integrated Quarticity Estimation

2011-24: Cristina Amado and Timo Teräsvirta: Conditional Correlation Models of Autoregressive Conditional Heteroskedasticity with Nonstationary GARCH Equations

2011-25: Stephen T. Ziliak: Field Experiments in Economics: Comment on an article by Levitt and List

2011-26: Rasmus Tangsgaard Varneskov and Pierre Perron: Combining Long Memory and Level Shifts in Modeling and Forecasting of Persistent Time Series

2011-27: Anders Bredahl Kock and Timo Teräsvirta: Forecasting Macroeconomic Variables using Neural Network Models and Three Automated Model Selection Techniques

2011-28: Anders Bredahl Kock and Timo Teräsvirta: Forecasting performance of three automated modelling techniques during the economic crisis 2007-2009

2011-29: Yushu Li: Wavelet Based Outlier Correction for Power Controlled Turning Point Detection in Surveillance Systems

2011-30: Stefano Grassi and Tommaso Proietti: Stochastic trends and seasonality in economic time

2011-31: Rasmus Tangsgaard Varneskov: Generalized Flat-Top Realized Kernel Estimation of Ex-Post Variation of Asset

2011-32: Christian Bach: Conservatism in Corporate Valuation

2011-33: Adrian Pagan and Don Harding: Econometric Analysis and Prediction of Recurrent Events

2011-34: Lars Stentoft: American Option Pricing with Discrete and Continuous Time Models: An Empirical Comparison