Which GARCH Model for Option Valuation?

By Peter Christoffersen and Kris Jacobs

Online Appendix

Sample Index Returns

In order to give an idea of the differences in returns over the sample, Figure A.1 plots the S&P 500 daily log returns and daily absolute log returns from CRSP for the option sample period. The vertical line separates the Sample A and B periods. It can be seen that the patterns in returns are fairly similar over the two periods, and that there are occasional outliers in returns.

ML Estimation of the Price of Risk from Returns Only

It is well-known that it is difficult to estimate the parameter λ precisely from ML estimation on returns data. Even though λ is estimated significantly in Tables 4 and 5, its standard errors are quite large. In Figure A.2, we present an empirical analysis designed to investigate if this has consequences for option valuation. We take the parameters estimated for the leverage model in Table 5, and evaluate the $RMSE error for deviations of the risk-neutral leverage parameter $\theta^* = \theta + \lambda$ from its ML value. To ensure that the resulting set of parameter values is meaningful, we adjust the parameter β_1 to keep the persistence of the risk-neutral process $\beta_1 + \beta_2\ (1 + \theta^{*2})$ fixed at its ML estimate. This exercise is designed to investigate to what extent more precise estimation of λ could improve the valuation error. Figure A.2 indicates that the answer is somewhat sample-specific, with a larger potential improvement for Sample B than for Sample A.

Implementation of the NLS Estimation Procedures

This section describes a number of technical issues regarding the estimation procedures. The first issue concerns the NLS procedure described in Section 2.2 when multiple cross-sections of options are used in estimation. In this case, we need to use an updating rule to link volatility on different dates. The updating from h_t to h_{t+1} is done using the observed daily returns, R_t on the underlying asset, by substituting

$$z_{t-1} = \left(R_{t-1} - r - \lambda\sqrt{h_{t-1}} + \frac{1}{2}h_{t-1}\right)/\sqrt{h_{t-1}} = \left[\left(R_{t-1} - r + \frac{1}{2}h_{t-1}\right)/\sqrt{h_{t-1}}\right] - \lambda.$$

into

$$h_t = \beta_0 + \beta_1 h_{t-1} + \beta_2 h_{t-1} f(z_{t-1})$$

which yields an updating function that exclusively involves observables

$$h_t = \beta_0 + \beta_1 h_{t-1} + \beta_2 h_{t-1} f\left(\left[\left(R_{t-1} - r + \frac{1}{2}h_{t-1}\right)/\sqrt{h_{t-1}}\right] - \lambda\right)$$

It can therefore be seen that when we use multiple cross-sections of options data, we use joint information on returns and options prices, but the information on returns only enters the objective function indirectly. This approach is also followed by Heston and Nandi (2000) and Hsieh and Ritchken (2000) in a discrete-time environment.

Note also from the updating rule that there is an identification problem when estimating under the risk-neutral probability measure. For example, in the case of the *Leverage* model, the return dynamic is

$$R_t = r - \frac{1}{2}h_t + \sqrt{h_t}z_t^*, \quad z_t^* \sim N(0,1).$$

and the risk-neutral volatility process is

$$h_t = \beta_0 + \beta_1 h_{t-1} + \beta_2 h_{t-1}\left(z_{t-1}^* - \theta - \lambda\right)^2$$

Notice that we cannot separately identify λ and θ from the two above equations as we can only estimate $\theta^* \equiv \theta + \lambda$. Notice that for the leverage model the volatility updating formula becomes

$$h_t = \beta_0 + \beta_1 h_{t-1} + \beta_2 h_{t-1}\left(\left[\left(R_{t-1} - r + \frac{1}{2}h_{t-1}\right)/\sqrt{h_{t-1}}\right] - (\lambda + \theta)\right)^2$$

So clearly we cannot separately identify λ and θ when estimating under the risk-neutral probability measure, as we only estimate $\theta^* \equiv \theta + \lambda$. Alternatively one can assume that $\lambda = 0$, which trivially identifies θ (see Hsieh and Ritchken 2000).

This identification problem is not a problem for option valuation, because all we need is the combination $\theta + \lambda$, but it is an important difference compared to the ML estimation using returns, where we separately identify λ and θ using the physical probability measure and then use the sum of the two parameters for option valuation. This remark applies for all of the models we investigate, as they all involve the same functional form in θ. It is important to remember that this issue is simply a consequence of the way in which the parameter λ enters the innovation function, and therefore a consequence of the assumptions on the discount factor (or the price of risk). Also, it must be emphasized that in principle one could identify λ and θ separately, for instance by specifying a joint likelihood function for returns and options prices. By using the updating rule returns data are only implicitly taken into account in the objective function, and this does not identify λ.

A second implementation issue that deserves elaboration is the treatment of initial conditional volatility. Bakshi, Cao and Chen (1997) treat the initial conditional volatility as an extra parameter in the implementation of a continuous-time option valuation model. This approach could also be followed for the estimation of discrete-time processes, but instead we follow Heston and Nandi (2000) and set the initial volatility equal to the unconditional volatility 250 days before the first Wednesday that is included in the sample. Subsequently, we use the volatility updating rule for the next 250 days to obtain the conditional volatility for the first Wednesday used in the sample.

A third issue regarding implementation is the size, nature and length of the option valuation sample. Our motivation in selecting the sample is essentially to use the longest time span possible, while limiting the number of contracts to limit computation time. There are several good reasons for preferring a long sample. First, as emphasized by Heston and Nandi (2000), if the option valuation model is a reliable one, it should price options well over any time horizon, because it has implications for the returns process as well as for option prices. Second, from a practical perspective, a long sample of option prices is preferable for the same reason a long sample of returns is preferable, because the time variation in prices and volatility will help identify the model's parameters. We use a long, three-year sample of options data to estimate the model's parameters. Extensive sensitivity analysis demonstrated that the results for this approach are robust. As mentioned in Section 3, we only include Wednesdays in the sample following the setup in Dumas, Fleming and Whaley (1998) and Heston and Nandi (2000). It would be prohibitively expensive in terms of computer time to use all trading days in the three-year sample. The design of our empirical experiment can therefore be summarized as follows: we believe that for estimation purposes, it is preferable to use 156 consecutive Wednesdays rather than 156 consecutive trading days.

Finally, the implementation of the Monte Carlo simulation is critically important. In order to compute the option prices numerically, we use 1000 simulated paths, and we use the empirical martingale method of Duan and Simonato (1999) to increase numerical efficiency. We also use stratified random numbers, antithetic variates, and a control variate technique using the Black-Scholes price as the control. To verify that 1000 draws are adequate, we repeated our analysis for a limited number of cases using 5000 draws and obtained identical results.

Exploring the Results from Option-Based Estimation

Figures A.3 and A.4 provide more insight into the valuation differences between the models by comparing the option prices from each model with the Black-Scholes price for the same option. For each model, the risk-neutral parameter estimates in Table 6 are used. The call price from each model is plotted against moneyness for three maturities. Similar figures are shown in Heston (1993) for his stochastic volatility model. The initial conditional volatility in each GARCH model is set to its unconditional value as implied by the assumed parameters, and the volatility used for Black-Scholes valuation is set to this conditional volatility. It must be noted that this setup ensures that the Black-Scholes price has the best possible chance of matching the GARCH prices, because the initial volatilities are the same across models. We see that the *Simple* model yields very small deviations from the Black-Scholes model across maturities. The other models, particularly the $Box - Cox$ model, display systematic differences from Black-Scholes, enabling the models to fit observed patterns in the data such as implied volatility smirks. The differences between the GARCH models and Black-Scholes are smaller for short maturities, partly as a result of the lower values of these options.

Table A.1 further elaborates on these findings by presenting $RMSEs by moneyness and maturity for the four models (for Sample A). A first interesting finding from these tables is that in certain cells there are notable differences between the *Leverage*, *News* and *Power* models. The differences between these three models are perhaps surprising, because on the basis of the overall fit in Table 7, one might have concluded that the three models yield nearly identical option prices. Table A.1 indicates that this is not the case.

It can also be seen from Table A.1 that even though the overall fit of the *Power&News* model is better than that of the other three models, this does not mean that this model does a better job of valuing all options. It is not a surprise that the most important improvements over the other three models are made for short and medium maturity options. The reason for this is that differences between GARCH models are more likely to be significant for short and medium-horizon forecasts, while they even out over long horizons where the conditional volatility approaches the unconditional. Furthermore, for the short and medium maturities improvements are made primarily for options that are in the money. The reason for this is that in-the-money options are more expensive and therefore carry more weight in the objective function.

Table A.2 breaks down out-of-sample model performance across moneyness and maturity. By simply comparing the two models with the best and worst overall out-of-sample fit, it becomes clear that comparing the models' valuation performance is less than straightforward. While the overall fit of the *Power&News* model is significantly worse than that of the *Leverage* model ($1.1571 versus $0.9777 in $RMSE terms), it still outperforms the *Leverage* model for some deep in-the-money options. However, for long-maturity in-the-money options, the performance of the *Power&News* model is relatively poor.

Calculating the Semiparametric Densities

In computing the option-implied risk neutral densities in Figure 5 we follow the semiparametric approach in Ait-Sahalia and Lo (1998). Their method consists of the following steps:

1) Calculate the implied Black-Scholes volatility, σ_i for each option in the sample. We use Sample A.

2) Estimate the three-dimensional kernel regression for implied volatility by

$$\hat{\sigma}\left(F_{t,\tau}, K, \tau\right) = \frac{\sum_{i=1}^{n} k_F\left(\frac{F_{t,\tau} - F_{t_i,\tau_i}}{h_F}\right) k_K\left(\frac{K_{t,\tau} - K_i}{h_K}\right) k_\tau\left(\frac{\tau - \tau_i}{h_\tau}\right)\sigma_i}{\sum_{i=1}^{n} k_F\left(\frac{F_{t,\tau} - F_{t_i,\tau_i}}{h_F}\right) k_K\left(\frac{K_{t,\tau} - K_i}{h_K}\right) k_\tau\left(\frac{\tau - \tau_i}{h_\tau}\right)}$$

where $F_{t,\tau} = S_t \exp\left(r_{t,\tau}\right)$ denotes the forward price on day t for maturity τ. We use the normal kernel and a bandwidth which is proportional to the standard deviation of each variable. The bandwidth is chosen to achieve an appropriate degree of smoothness in the risk-neutral densities below.

3) The call pricing function is estimated using the Black-Scholes functional form as in

$$\hat{C}\left(S_t, K, \tau, r_{t,\tau}\right) = C_{BS}\left(F_{t,\tau}, K, \tau, r_{t,\tau}; \hat{\sigma}\left(F_{t,\tau}, K, \tau\right)\right)$$

4) The risk neutral density for the spot price on maturity date, $T = t + \tau$, is calculated as a simple function of the second derivative of the semiparametric option price with respect to the strike price:

$$\hat{f}_t^*\left(S_T\right) = \exp\left(r_{t,\tau}\right)\left[\frac{\partial^2 C_{BS}\left(F_{t,\tau}, K, \tau, r_{t,\tau}; \hat{\sigma}\left(F_{t,\tau}, K, \tau\right)\right)}{\partial K^2}\right]_{|K=S_T}$$

We calculate this derivative numerically across a grid of strike prices for each horizon setting the current spot price and interest rate equal to their average sample values.

5) In order to plot the density against log returns rather than future spot prices, we use the transformation

$$\frac{\partial}{\partial u} \Pr\left(\ln\left(\frac{S_T}{S_t}\right) \leq u\right) = S_t \exp(u) \hat{f}_t^* \left(S_t \exp(u)\right)$$

When calculating the risk neutral densities implied by the GARCH MLs in Table 5 we simply replace the implied volatilities from the observed option prices with the implied volatilities from the fitted option prices in the GARCH option pricing model and redo steps 1-5 above.

Figure A.1: S&P500 Returns Across Option Sample Period

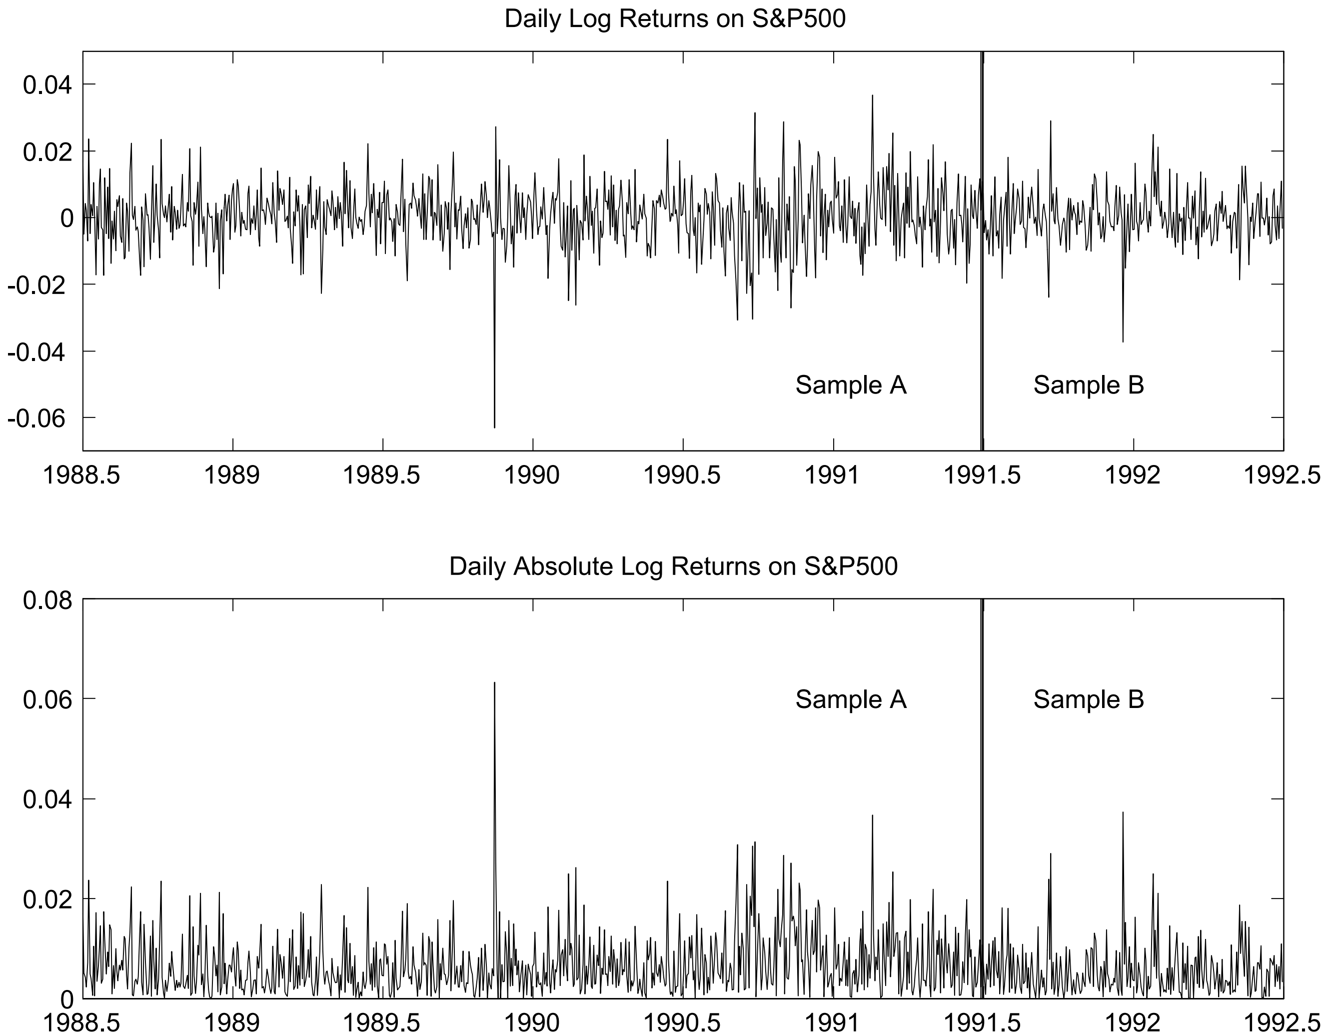

Notes: The top panel shows the daily log returns on the S&P500 index. The data is shown from June 1, 1988 through May 31, 1992. The vertical line delimits Sample A, which corresponds to June 1, 1988 through May 31, 1991, from Sample B, which corresponds to June 1, 1991 through May 31, 1992. The bottom panel shows the daily absolute returns on the S&P500 index for the same periods.

Figure A.2: Option Valuation \$RMSE: Perturbing the MLE of the Leverage Effect

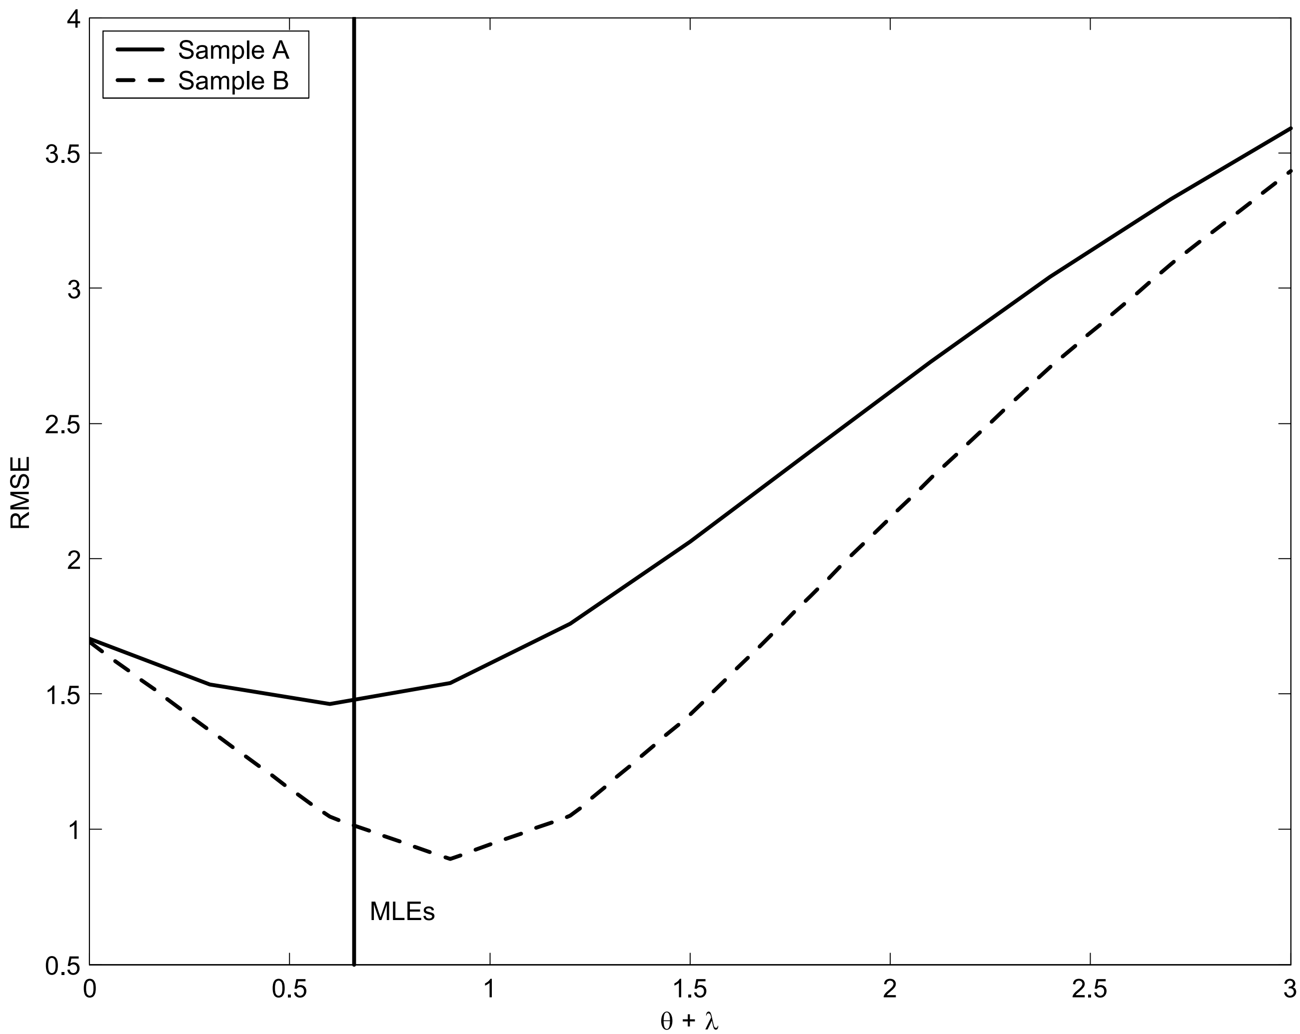

Notes: The return-based ML estimates of the *Leverage* model from Table 5 are used to calculate the \$RMSE of the options in Sample A and B. The risk neutral leverage parameter $\theta^* = \theta + \lambda$ is perturbed from its ML value while the persistence parameter β_1 is adjusted correspondingly to keep fixed the persistence of the risk neutral process, $\beta_1 + \beta_2 \left(1 + \theta^{*2}\right)$. The horizontal axis denotes the range of values for θ^* and the vertical axis denotes the \$RMSE of the options. The vertical line shows the MLE estimate of θ^*. The solid graph shows the \$RMSE from option Sample A and the dashed line from option Sample B.

Figure A.3: Model Price Less Black-Scholes Price. Various Models

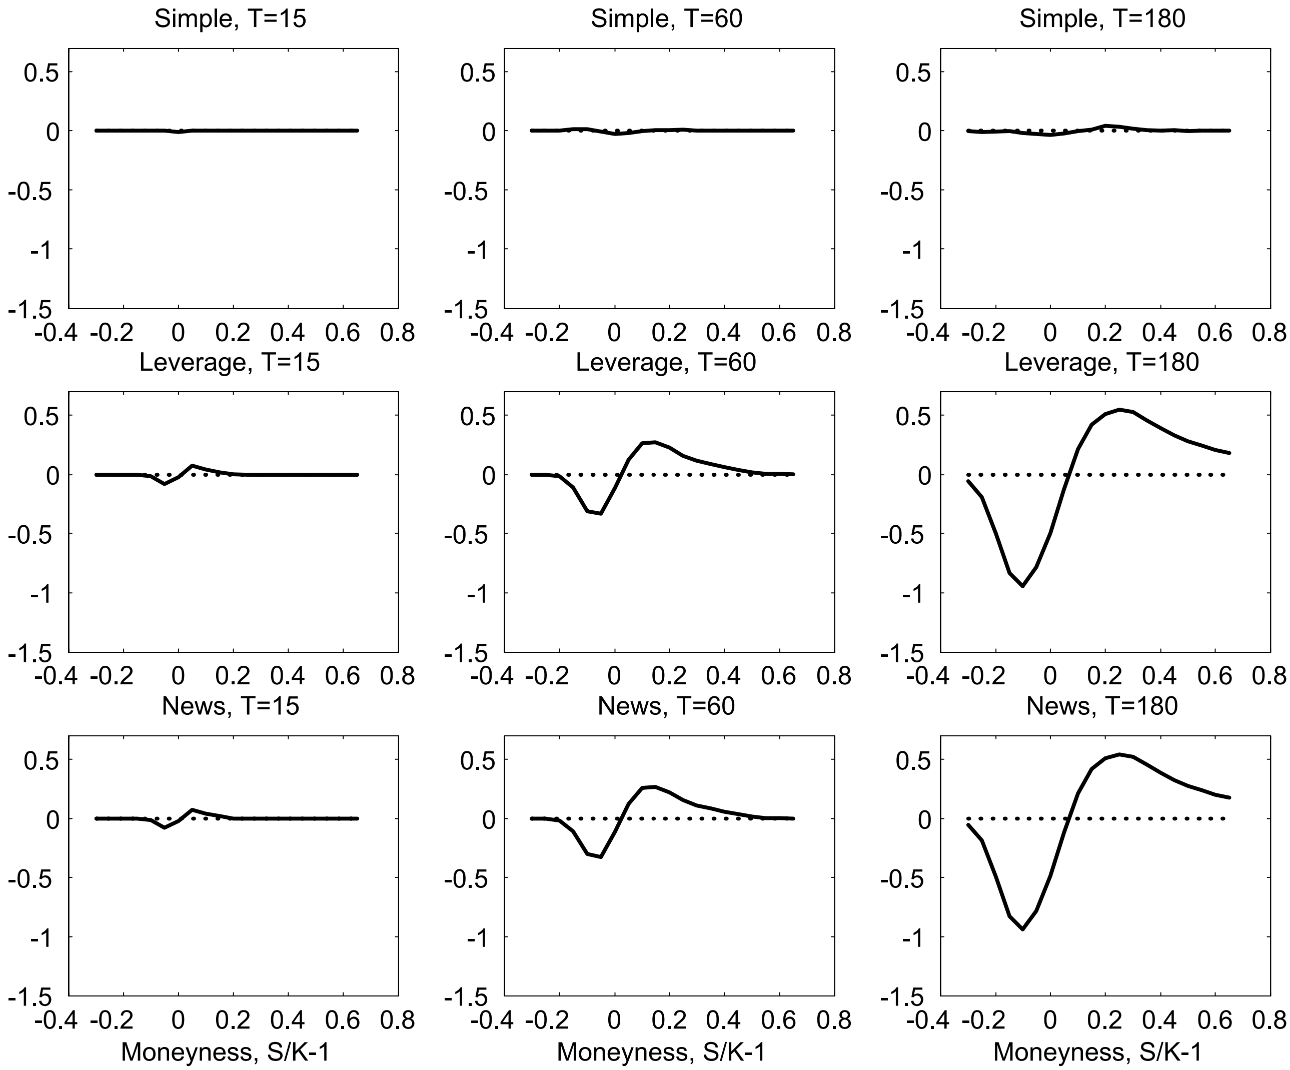

Notes: Each line shows the option price from a particular GARCH model less the Black-Scholes price of the same option. Horizontal dashed lines are plotted at zero. The horizontal axis denotes the moneyness of the option, $S/K-1$. The parameters of each GARCH model are set equal to the NLS estimates from Table 6. The initial conditional variance in the GARCH model and the volatility in the Black-Scholes model are both set equal to the unconditional value implied by the parameters in the GARCH model. The columns of plots correspond to 15, 60 and 180 days to maturity respectively. Each row of plots corresponds to a particular GARCH model.

Figure A.4: Model Price Less Black-Scholes Price. Various Models (cont.)

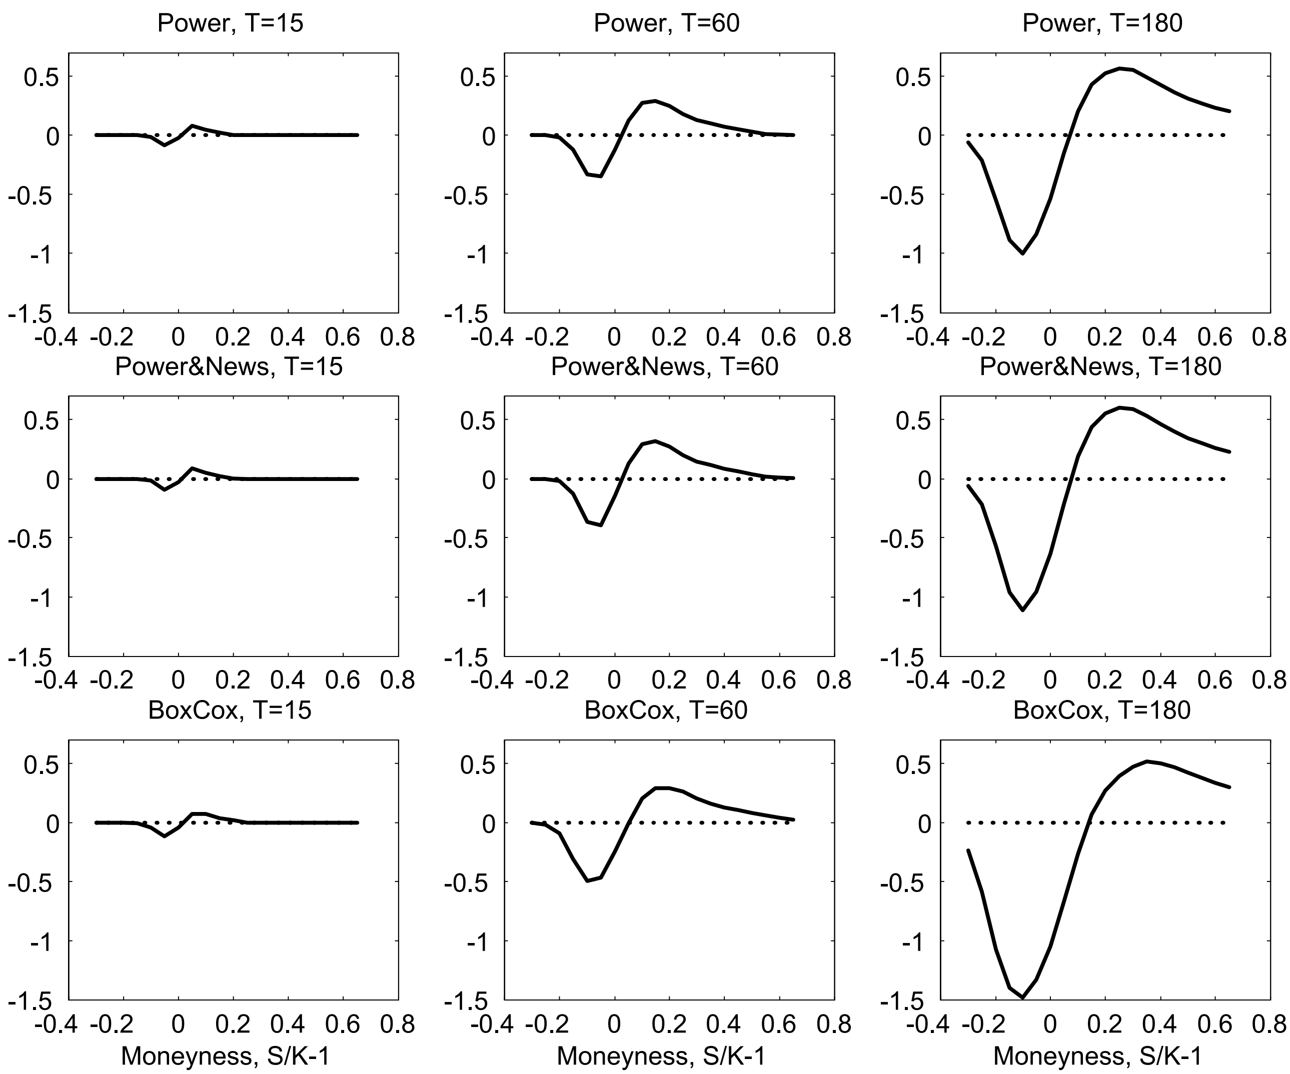

Notes: Each line shows the option price from a particular GARCH model less the Black-Scholes price of the same option. Horizontal dashed lines are plotted at zero. The horizontal axis denotes the moneyness of the option, $S/K-1$. The parameters of each GARCH model are set equal to the NLS estimates from Table 6. The initial conditional variance in the GARCH model and the volatility in the Black-Scholes model are both set equal to the unconditional value implied by the parameters in the GARCH model. The columns of plots correspond to 15, 60 and 180 days to maturity respectively. Each row of plots corresponds to a particular GARCH model.

Table A.1: $RMSE on Sample A from NLS Estimates on Sample A. By Moneyness and Maturity

Simple Model

	DTM<60	60<DTM<180	180<DTM
S/X < .94	0.9616	1.5246	2.2977
.94 < S/X < .97	0.9841	1.5720	2.0373
.97 < S/X < 1.00	0.9155	1.5963	1.9129
1.00 < S/X < 1.03	0.9991	1.8677	2.1563
1.03 < S/X < 1.06	1.2191	2.1494	1.7587
1.06 < S/X	1.0448	1.8351	2.1288

Leverage Model

	DTM<60	60<DTM<180	180<DTM
S/X < .94	0.6470	0.9089	1.3186
.94 < S/X < .97	0.7692	0.9915	1.3204
.97 < S/X < 1.00	0.8277	1.0415	1.3636
1.00 < S/X < 1.03	0.8304	1.1628	1.3381
1.03 < S/X < 1.06	0.8701	1.2720	1.2241
1.06 < S/X	0.8033	1.1245	1.3125

News Model

	DTM<60	60<DTM<180	180<DTM
S/X < .94	0.6531	0.9123	1.3120
.94 < S/X < .97	0.7705	0.9917	1.3126
.97 < S/X < 1.00	0.8236	1.0383	1.3575
1.00 < S/X < 1.03	0.8221	1.1579	1.3354
1.03 < S/X < 1.06	0.8635	1.2674	1.2168
1.06 < S/X	0.8004	1.1218	1.3093

Power Model

	DTM<60	60<DTM<180	180<DTM
S/X < .94	0.6281	0.8915	1.2989
.94 < S/X < .97	0.7477	0.9878	1.3223
.97 < S/X < 1.00	0.8174	1.0478	1.3682
1.00 < S/X < 1.03	0.8329	1.1715	1.3419
1.03 < S/X < 1.06	0.8784	1.2642	1.2481
1.06 < S/X	0.8076	1.1059	1.3169

News&Power Model

	DTM<60	60<DTM<180	180<DTM
S/X < .94	0.6322	0.9033	1.2228
.94 < S/X < .97	0.7377	1.0099	1.2986
.97 < S/X < 1.00	0.7993	1.0601	1.3427
1.00 < S/X < 1.03	0.7839	1.1573	1.3421
1.03 < S/X < 1.06	0.8193	1.1995	1.2692
1.06 < S/X	0.7640	1.0129	1.3344

Box-Cox Model

	DTM<60	60<DTM<180	180<DTM
S/X < .94	0.6446	0.9095	1.3184
.94 < S/X < .97	0.7659	0.9951	1.3193
.97 < S/X < 1.00	0.8254	1.0441	1.3629
1.00 < S/X < 1.03	0.8303	1.1648	1.3343
1.03 < S/X < 1.06	0.8723	1.2716	1.2233
1.06 < S/X	0.8052	1.1233	1.3086

Notes: We report $RMSE from various GARCH options valuation models on Sample A, which denotes June 1, 1988 - May 31, 1991. The $RMSE is reported for three maturity bins and six moneyness bins. The parameters in the GARCH option valuation models are estimated minimizing $MSE on Sample A itself.

Table A.2: $RMSE on Sample B from NLS Estimates on Sample A. By Moneyness and Maturity

Simple Model

	DTM<60	60<DTM<180	180<DTM
S/X < .94	0.9973	1.9716	2.7712
.94 < S/X < .97	1.1416	1.8755	1.8108
.97 < S/X < 1.00	0.9931	1.3756	1.1216
1.00 < S/X < 1.03	0.5711	1.1169	1.3577
1.03 < S/X < 1.06	0.7546	1.4163	1.9170
1.06 < S/X	0.9460	1.8209	2.3680

Leverage Model

	DTM<60	60<DTM<180	180<DTM
S/X < .94	0.5352	0.9331	1.4045
.94 < S/X < .97	0.7836	1.1320	1.3652
.97 < S/X < 1.00	0.8880	1.1886	1.1978
1.00 < S/X < 1.03	0.7385	0.9870	1.0201
1.03 < S/X < 1.06	0.6404	0.9713	0.9516
1.06 < S/X	0.8008	1.1000	0.8881

News Model

	DTM<60	60<DTM<180	180<DTM
S/X < .94	0.5458	0.9376	1.3806
.94 < S/X < .97	0.7941	1.1359	1.3354
.97 < S/X < 1.00	0.8978	1.1883	1.1797
1.00 < S/X < 1.03	0.7395	0.9830	0.9948
1.03 < S/X < 1.06	0.6353	0.9657	0.9389
1.06 < S/X	0.7979	1.0971	0.8786

Power Model

	DTM<60	60<DTM<180	180<DTM
S/X < .94	0.5498	0.9624	1.4468
.94 < S/X < .97	0.8083	1.1729	1.4406
.97 < S/X < 1.00	0.9241	1.2323	1.2659
1.00 < S/X < 1.03	0.7595	1.0090	1.1040
1.03 < S/X < 1.06	0.6327	0.9609	1.0110
1.06 < S/X	0.7849	1.0523	0.9251

News&Power Model

	DTM<60	60<DTM<180	180<DTM
S/X < .94	0.6792	1.0682	1.3759
.94 < S/X < .97	0.9694	1.3163	1.4072
.97 < S/X < 1.00	1.1400	1.3925	1.3475
1.00 < S/X < 1.03	0.9278	1.1518	1.1585
1.03 < S/X < 1.06	0.6577	1.0184	1.1197
1.06 < S/X	0.7260	0.9629	1.0038

Box-Cox Model

	DTM<60	60<DTM<180	180<DTM
S/X < .94	0.5411	0.9414	1.4098
.94 < S/X < .97	0.7894	1.1407	1.3742
.97 < S/X < 1.00	0.8941	1.1947	1.2105
1.00 < S/X < 1.03	0.7396	0.9868	1.0272
1.03 < S/X < 1.06	0.6373	0.9654	0.9583
1.06 < S/X	0.7984	1.0921	0.8886

Notes: We report $RMSE from various GARCH options valuation models on Sample B, which denotes June 1, 1991 - May 31, 1992. The $RMSE is reported for three maturity bins and six moneyness bins. The parameters in the GARCH option valuation models are estimated minimizing $MSE on Sample A, which denotes June 1, 1988 - May 31, 1991.