

THE INFORMATION CONTENT OF IMPLIED VOLATILITY IN AGRICULTURAL COMMODITY MARKETS

Pierre Giot¹

May 2002

Abstract

In this paper we compare the incremental information content of lagged implied volatility to GARCH models of conditional volatility for a collection of agricultural commodities traded on the New York Board of Trade. We also assess the relevance of the additional information provided by the implied volatility in a risk management framework. It is first shown that past squared returns only marginally improve the information content provided by the lagged implied volatility. Secondly, Value-at-Risk (VaR) models that rely exclusively on lagged implied volatility perform as well as VaR models where the conditional variance is modelled according to GARCH type processes. These results indicate that the implied volatility for options on future contracts in agricultural commodity markets has a high information content regarding conditional variance and VaR forecasts.

Keywords: Implied volatility, GARCH, Value-at-Risk, futures, agricultural commodity markets

JEL classification: C52, C53, G15, Q13

¹Department of Business Administration, University of Namur, Rempart de la Vierge 8, B-5000 Namur, Belgium, Tel.: +3281724887 and Center for Operations Research and Econometrics, UCL, Belgium; email: pierre.giot@fundp.ac.be

While remaining responsible for any errors in this paper, the author would like to thank Sébastien Laurent and Paul Reding for useful remarks and suggestions. We also thank the New York Board of Trade for the availability of the price and implied volatility time series on their web site.

1 Introduction

In an option pricing framework, the implied volatility σ_{imp}^2 is defined as the volatility that, when plugged in the option pricing formula, equates the theoretical price of the option with the observed market price. For example, the Black and Scholes (BS, 1973) implied volatility is computed by ‘inverting’ the BS option pricing formula such that σ_{imp}^2 can be determined as a function of the market price and characteristics of the option. While volatility is assumed constant in the BS framework, most empirical studies show that implied volatility is strike (smile and skew effects) and maturity (term structure of implied volatility) dependent. Because implied volatility reflects the average expected (by market participants) volatility over the life of the option, most market practitioners consider this measure of volatility as the most trustworthy forecast of the short-term volatility of the underlying asset. Further information is available in Hull (2000) or Alexander (2001) for example.

In a volatility forecasting setting of the ARCH¹ type, a number of recent empirical studies look at the relevance of the additional information provided by the lagged implied volatility and assess how it improves on the information given by the past squared returns when it is included in the conditional variance equation. Examples are Day and Lewis (1992), who compare the information content of implied volatility of call options on the S&P100 index to GARCH type conditional volatility; Xu and Taylor (1995) who focus on the informational efficiency of the PHLX currency options market; Blair, Poon, and Taylor (2001) who compare the information content of implied volatilities and intraday returns when short-term index volatility is to be forecasted.

In this paper we deal with the same topic for a collection of agricultural commodities traded on the New York Board of Trade (NYBOT) but we also extend the previous analysis by assessing the relevance of the additional information provided by the implied volatility in a risk management framework of the Value-at-Risk type, i.e. when daily VaR levels are to be forecasted. Our empirical application focuses on several agricultural commodity products (cacao, coffee and sugar future contracts) for which implied volatility computed from short-term options on nearby futures is readily available. More precisely we estimate conditional volatility models which include past squared returns (ARCH type models) and lagged implied volatility (options information only) and conclude that past squared returns only marginally improve the information content provided by the lagged implied volatility. In a second stage, we show that VaR models that rely exclusively on lagged implied volatility perform as well as VaR models where the conditional variance is modelled according to ARCH type processes. Our results thus show that the implied volatility for options on future contracts in cocoa, coffee and sugar commodity markets has a high information content regarding conditional variance and VaR forecasts.

¹See Engle (1995) for a review of ARCH models.

The rest of the paper is organized in the following way. In Section 2, we present the conditional volatility models that take as inputs the squared returns and/or the implied volatility. The empirical application for the cacao, coffee and sugar nearby future contracts is given in Section 3. Section 4 concludes.

2 Conditional volatility and the information content of implied volatility

For a given commodity, we consider a collection of daily returns for future prices, $r_t = \ln(F_t) - \ln(F_{t-1})$, with $t = 1 \dots T$ and F_t is the (closing) price of the future contract on day t . Allowing for possible autocorrelation in the returns, their dynamics is characterized by an AR(p) process:

$$r_t = \rho_0 + \rho_1 r_{t-1} + \dots + \rho_p r_{t-p} + e_t \quad (1)$$

with $e_t = \sqrt{h_t} \epsilon_t$. Conditional heteroskedasticity for the error term e_t is modelled by a skewed Student GARCH(1,1) model.² Note that we focus directly on a density distribution that allows skewness and excess kurtosis as previous empirical work has shown at length that asset returns (at least on a daily basis) are almost always leptokurtic and often exhibit a non-zero skewness. Other candidates for the density distribution of the error term would be the normal distribution (whose performance is poor), the Student distribution (acceptable performance, but does not allow for skewness) or the mixture of normal distributions (another interesting candidate). See for example Alexander (2001), Mittnik and Paoletta (2000) or Giot and Laurent (2001) who compare the volatility forecasting performance of the normal, Student and skewed Student APARCH model.

To assess the information content of the lagged implied volatility of traded options for the underlying future contracts, the model is estimated without and with the lagged implied volatility, and with the lagged implied volatility only (i.e. no GARCH effects in this specification as this third model only uses options information). One thus estimates a model for the market returns r_t where the specification of the conditional volatility is successively given by

$$h_t = \omega + \alpha_1 e_{t-1}^2 + \beta_1 h_{t-1}, \quad (3)$$

²According to Lambert and Laurent (2001), the innovation process ϵ is said to be (standardized) skewed Student distributed if:

$$f(\epsilon|\xi, v) = \begin{cases} \frac{2}{\xi + \frac{1}{\xi}} sg[\xi(s\epsilon + m)|v] & \text{if } \epsilon < -\frac{m}{s} \\ \frac{2}{\xi + \frac{1}{\xi}} sg[(s\epsilon + m)/\xi|v] & \text{if } \epsilon \geq -\frac{m}{s} \end{cases} \quad (2)$$

where $g(\cdot|v)$ is the symmetric (unit variance) Student density and ξ is the asymmetry coefficient;³ m and s^2 are respectively the mean and the variance of the non-standardized skewed Student. The GARCH(1,1) model was put forward by Bollerslev (1986).

$$h_t = \omega + \alpha_1 e_{t-1}^2 + \beta_1 h_{t-1} + \eta \sigma_{imp,t-1}^2 \quad (4)$$

and

$$h_t = \omega + \eta \sigma_{imp,t-1}^2. \quad (5)$$

The distribution of the error term ϵ_t is as given by Equation (2). Note that Equation (5) is nested within Equation (4) as it only includes options market volatility in the conditional variance specification, while coefficient η in Equation (4) can be interpreted as a measure of the incremental information of lagged implied volatility with respect to the information provided by lagged squared returns.

3 Empirical application and Value-at-Risk forecasts

3.1 Conditional volatility and lagged implied volatility

Our empirical application deals with three agricultural commodities (cacao, coffee and sugar future contracts) traded on the NYBOT. Trading in future contracts on this exchange is very active with multiple delivery dates for each commodity. We focus on the analysis of the nearby future contracts for which implied volatility is directly available.⁴ Descriptive characteristics for the returns series are given in Table 1, while we report the price paths of the nearby future contracts and the corresponding time series of the implied volatilities in Figures 1-3. Over our estimation sample, the returns of the nearby future contracts for the three commodities were leptokurtic, positively (for the cocoa and coffee) or negatively (for the sugar) skewed and exhibited heteroskedasticity.

Estimation results⁵ for the 3/1/1994 - 30/12/1999 period are given in Table 2. In the GARCH specification, coefficient β_1 is close to 1 which is consistent with the observed clustering of volatility. As in previous studies focusing on financial assets (see the references given in the introduction), the lagged implied volatility takes up most of the GARCH effect when included in the model as evidenced by the sharp drop in the value of coefficient β_1 and the significant coefficient η . The distribution of returns features fat tails as ν is close to 7 for the cocoa future contracts, and close to 4 for the coffee and sugar contracts.

Table 3 gives the maximum log-likelihoods after ML estimation for the three commodities and for the three possible specifications. In this table, the second and third rows are of particular

⁴The time series of implied volatility is provided by the NYBOT and is based on the average nearby in-the-money, at-the-money and out-of-the-money call and put options.

⁵We report results for coefficients β_1 and η only as these coefficients are the most important for our analysis. The lag structure in the AR(p) specification is set to 3, and a (1,1) structure is chosen for the GARCH process. This ensures that residuals and squared standardized residuals are not correlated.

interest as they allow a comparison of the nested specifications given by Equations (4) and (5). For the cocoa, coffee and sugar nearby future prices, the decrease in maximum log-likelihoods when switching from Equation (4) to Equation (5) is equal to 0.55, 4.8 and 6.09 respectively. Using twice those values and with the critical value for the χ^2 with 2 degrees of freedom equal to 5.99 at the five percent level, the null hypothesis that past squared returns (i.e. GARCH effects) add no significant volatility information in addition to the lagged implied volatility is not rejected for the cocoa future contracts (thus the options market is informationally efficient for the cocoa contracts), while it is rejected for the coffee and sugar future contracts. Nevertheless the differences in the maximum log-likelihoods are rather small and, at the one percent level with the critical value now at 9.21, the previous hypothesis is thus barely rejected for the coffee future contracts.

3.2 Short-term VaR forecasts

The 1996 Amendment to the 1988 Basel Accord for market risk put forward a new approach as to how the market risk capital requirement should be computed, allowing the use of an internal model to compute the maximum loss over 10 trading days at a 99% confidence level. This set the stage for the Value-at-Risk models, which can be broadly defined as quantitative tools whose goal is to assess the possible loss that can be incurred by a financial institution over a given time period and for a given portfolio of assets.⁶ In our framework, the asset is the agricultural commodity (cacao, coffee or sugar) whose relevant market price is the nearby future price. We characterize daily market risk for the nearby future prices, our time horizon is thus set to one day and we consider five confidence levels ranging from 95% to 99.75%.

Moreover, the one-day VaR level is computed for commodity traders having either bought the future contract (long position) or short-sold it (short position).⁷ In the first case, the risk comes from a drop in the price of the future contract, while the trader loses money when the price increases in the second case. Correspondingly, one focuses in the first case on the left side of the distribution of returns, and on the right side of the distribution in the second case. From a statistical point of view, the long and short VaR levels require the computation of the left and right quantiles as forecasted by the econometric model given by Equations (1), (2), (3), (4) and (5). More precisely, the VaRs for long and short positions are given by $st_{\alpha,v,\xi}h_t$ and $st_{1-\alpha,v,\xi}h_t$, with $st_{\alpha,v,\xi}$ being the left quantile at $\alpha\%$ for the skewed Student distribution with v degrees of freedom and asymmetry coefficient ξ ($st_{1-\alpha,v,\xi}$ is the corresponding right quantile).⁸ Assessment

⁶Further general information about VaR techniques and regulation issues are available in Dowd (1998), Jorion (2000) or Saunders (2000).

⁷An asset is short-sold by a trader when it is first borrowed and subsequently sold on the market. By doing this, the trader hopes that the price will fall, so that he can then buy the asset at a lower price and give it back to the lender.

⁸See Giot and Laurent (2001) for full analytical results and expressions for the quantile functions of the skewed

of the VaR performances of the three competing specifications requires the computation of each model's empirical failure rate.⁹ The statistical test $H_0 : f = \alpha$ against $H_1 : f \neq \alpha$, where f is the failure rate (estimated by \hat{f} , the empirical failure rate) is made using the Kupiec LR test (see Kupiec, 1995).

In-sample VaR estimation for the three specifications of the conditional variance, corresponding failure rates for long and short trading positions in the future contracts and outcomes of the Kupiec LR test are summarized in Table 4. Empirical failure rates which differ significantly from their theoretical values are highlighted in bold. Results given in this table indicate that all models perform remarkably well as there are almost no rejections. This table also indicates that the 'options information only' specification for the conditional variance yields equally good VaR results as the full GARCH (with lagged implied volatility) model. In other words, the volatility information impounded in the lagged implied volatility ensures adequate long and short VaR forecasts, without the need of using past squared returns.

Finally, full out-of-sample long and short VaR results for the coffee nearby future prices are given in Table 5. In this iterative procedure, the first estimation sample is the complete sample for which the data is available less the last four years (this latter period being the forecast sample). The models are estimated, the long and short VaRs are predicted and compared with the first return in the forecast sample. Thereafter, the estimation sample is augmented, the models are re-estimated and the VaRs are computed and compared with the second return in the forecast sample. We iterate the procedure until all days (less the last one) have been included in the estimation sample. These out-of-sample VaR results indicate that the three models perform adequately as all failure rates are statistically equal to their theoretical values (in this table a bold figure would indicate that the corresponding VaR is significantly different from the theoretical value, and there are none). These results are thus in agreement with the in-sample VaR forecasts.

4 Conclusion

In this paper we put forward an extension of the studies by Day and Lewis (1992) and Xu and Taylor (1995) and compared the incremental information content of lagged implied volatility to skewed Student GARCH models of conditional volatility for a collection of agricultural commodities (cacao, coffee and sugar nearby future contracts) traded on the New York Board of Trade. As in the previous studies on financial assets, it is shown that past squared returns only marginally improve the information content provided by the lagged implied volatility. In a second stage, we

⁹Empirical failure rates are computed by determining the proportion of demeaned returns e_{t+1} smaller (for the long positions) and larger (for the short positions) than the VaRs given by $st_{\alpha,v,\xi}h_t$ and $st_{1-\alpha,v,\xi}h_t$. If the VaR model is correctly specified, the failure rates should be equal to the pre-specified VaR levels.

also assessed the relevance of the additional information provided by the implied volatility in a Value-at-Risk forecasting framework. Our in-sample and out-of-sample analysis show that VaR models that rely exclusively on lagged implied volatility perform as well as VaR models where the conditional variance is modelled according to GARCH type processes. These results indicate that the implied volatility for options on future contracts has a high information content regarding conditional variance and VaR forecasts of the underlying future contracts.

References

- ALEXANDER, C. (2001): *Market models*. Wiley.
- BLACK, F., AND M. SCHOLES (1973): “The pricing of options and corporate liabilities,” *Journal of Political Economy*, 81, 637–659.
- BLAIR, B., S. POON, AND S. TAYLOR (2001): “Forecasting S&P100 volatility: the incremental information content of implied volatilities and high-frequency index returns,” *Journal of Econometrics*, 105, 5–26.
- BOLLERSLEV, T. (1986): “Generalized autoregressive conditional heteroskedasticity,” *Journal of Econometrics*, 31, 307–327.
- DAY, T., AND C. LEWIS (1992): “Stock market volatility and the information content of stock index options,” *Journal of Econometrics*, 52, 267–287.
- DOWD, K. (1998): *Beyond Value-at-Risk*. Wiley.
- ENGLE, R. (1995): *ARCH selected readings*. Oxford University Press, Oxford.
- GIOT, P., AND S. LAURENT (2001): “Value-at-Risk for long and short trading positions,” CORE DP 2001/22, Maastricht University METEOR RM/01/005.
- HULL, J. (2000): *Options, futures and other derivatives*. Prentice-Hall.
- JORION, P. (2000): *Value-at-Risk*. McGraw-Hill.
- KUPIEC, P. (1995): “Techniques for verifying the accuracy of risk measurement models,” *Journal of Derivatives*, 2, 173–84.
- LAMBERT, P., AND S. LAURENT (2001): “Modelling Financial Time Series Using GARCH-Type Models and a Skewed Student Density,” Mimeo, Université de Liège.
- MITTNIK, S., AND M. PAOLELLA (2000): “Conditional Density and Value-at-Risk Prediction of Asian Currency Exchange Rates,” *Journal of Forecasting*, 19, 313–333.

SAUNDERS, A. (2000): *Financial institutions management*. McGraw-Hill.

XU, X., AND S. TAYLOR (1995): “Conditional volatility and the informational efficiency of the PHLX currency options market,” *Journal of Banking and Finance*, 19, 803–821.

Table 1: Descriptive statistics

	Cocoa	Coffee	Sugar
Annual s.d.	26.62	50.10	31.07
Skewness	0.655	0.156	-0.878
Excess Kurtosis	3.302	8.233	8.456
Minimum	-5.316	-22.064	-17.115
Maximum	9.962	23.773	11.619
$Q^2(10)$	82.13	131.71	72.68

Descriptive statistics for the daily returns (expressed in %) of the corresponding commodity. All values are computed using PcGive. $Q^2(10)$ is the Ljung-Box Q-statistic of order 10 on the squared returns.

Table 2: Skewed Student GARCH without/with lagged implied volatility

	Cocoa		Coffee		Sugar	
β_1	0.964 (0.009)	0.433 (0.337)	0.879 (0.034)	0 (.)	0.964 (0.010)	0.432 (0.286)
η	-	0.444 (0.254)	-	0.974 (0.132)	-	0.552 (0.282)

Estimated coefficients and standard errors for the skewed Student GARCH model without (β_1 only) and with lagged implied volatility included in the conditional variance specification (β_1 and η). The estimation period is 3/1/1994 - 30/12/1999.

Table 3: Log-likelihoods

	Cocoa	Coffee	Sugar
st GARCH	-2778.98	-3611.36	-2878.62
st GARCH with $\sigma_{imp,t-1}^2$	-2773.39	-3599.93	-2863.84
$\sigma_{imp,t-1}^2$ only	-2773.94	-3604.73	-2869.93

Log-likelihoods after ML estimation for the skewed Student GARCH(1,1) model without and with the lagged implied volatility, and with the lagged implied volatility only (i.e. no GARCH effects in this third specification). The estimation period is 3/1/1994 - 30/12/1999.

Table 4: In-sample VaR results for the cacao, coffee and sugar nearby future prices

	VaR for long positions					VaR for short positions				
	5%	2.5%	1%	0.5%	0.25%	5%	2.5%	1%	0.5%	0.25%
Cacao nearby future prices										
st GARCH	5.01	2.74	0.73	0.33	0.20	5.15	2.67	1.00	0.53	0.33
st GARCH with $\sigma_{imp,t-1}^2$	5.35	2.81	0.67	0.47	0.27	4.68	2.74	0.87	0.53	0.33
$\sigma_{imp,t-1}^2$ only	5.35	2.81	0.67	0.40	0.27	4.68	2.74	0.87	0.53	0.33
Coffee nearby future prices										
st GARCH	5.28	2.27	1.07	0.53	0.27	5.55	2.07	0.73	0.33	0.27
st GARCH with $\sigma_{imp,t-1}^2$	5.08	2.61	1.20	0.40	0.20	5.21	2.54	0.60	0.40	0.20
$\sigma_{imp,t-1}^2$ only	5.21	2.47	1.20	0.47	0.20	4.95	2.81	0.67	0.53	0.20
Sugar nearby future prices										
st GARCH	5.01	2.81	1.14	0.80	0.40	5.28	2.34	0.67	0.20	0
st GARCH with $\sigma_{imp,t-1}^2$	5.21	2.87	1.27	0.67	0.40	5.41	2.87	0.47	0.27	0
$\sigma_{imp,t-1}^2$ only	5.28	2.94	1.27	0.60	0.47	5.55	2.67	0.40	0.27	0

Failure rates for the skewed Student GARCH(1,1) model without and with the lagged implied volatility, and with the lagged implied volatility only (i.e. no GARCH effects in this third specification). A bold figure indicates that the corresponding VaR is significantly different (LR test) from the theoretical value. In-sample VaR results for the 3/1/1994 - 30/12/1999 period.

Table 5: Out-of-sample VaR results for the coffee nearby future prices

	VaR for long positions					VaR for short positions				
	5%	2.5%	1%	0.5%	0.25%	5%	2.5%	1%	0.5%	0.25%
st GARCH	5.46	2.78	1.39	0.69	0.40	5.95	2.38	0.79	0.30	0.30
st GARCH with $\sigma_{imp,t-1}^2$	5.75	3.17	1.59	0.79	0.30	5.85	2.48	0.59	0.40	0.20
$\sigma_{imp,t-1}^2$ only	6.15	2.98	1.59	0.59	0.30	5.36	2.68	0.89	0.50	0.10

Failure rates for the skewed Student GARCH(1,1) model without and with the lagged implied volatility, and with the lagged implied volatility only (i.e. no GARCH effects in this third specification). A bold figure would indicate that the corresponding VaR is significantly different (LR test) from the theoretical value. Out-of-sample VaR results for four years of data.

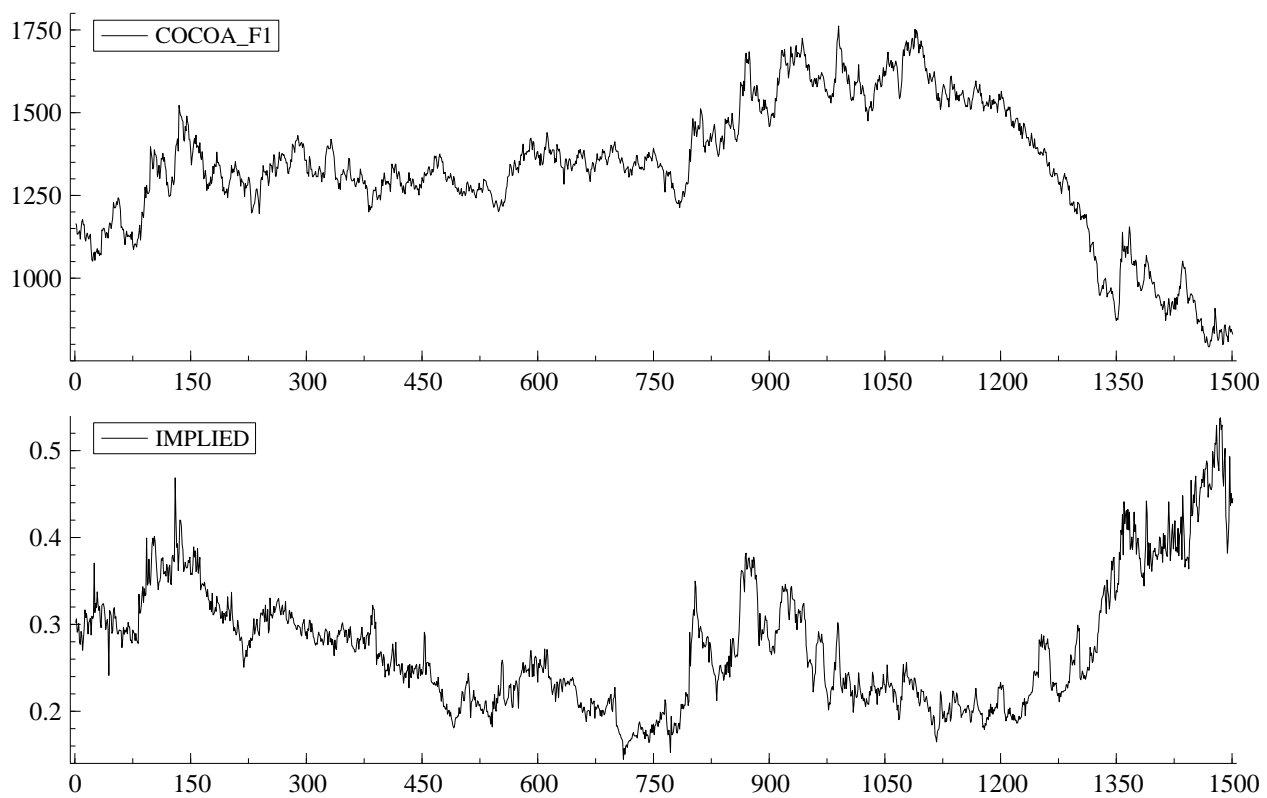


Figure 1: Nearby future cocoa daily prices and implied volatility. The time period is 3/1/1994 - 30/12/1999.

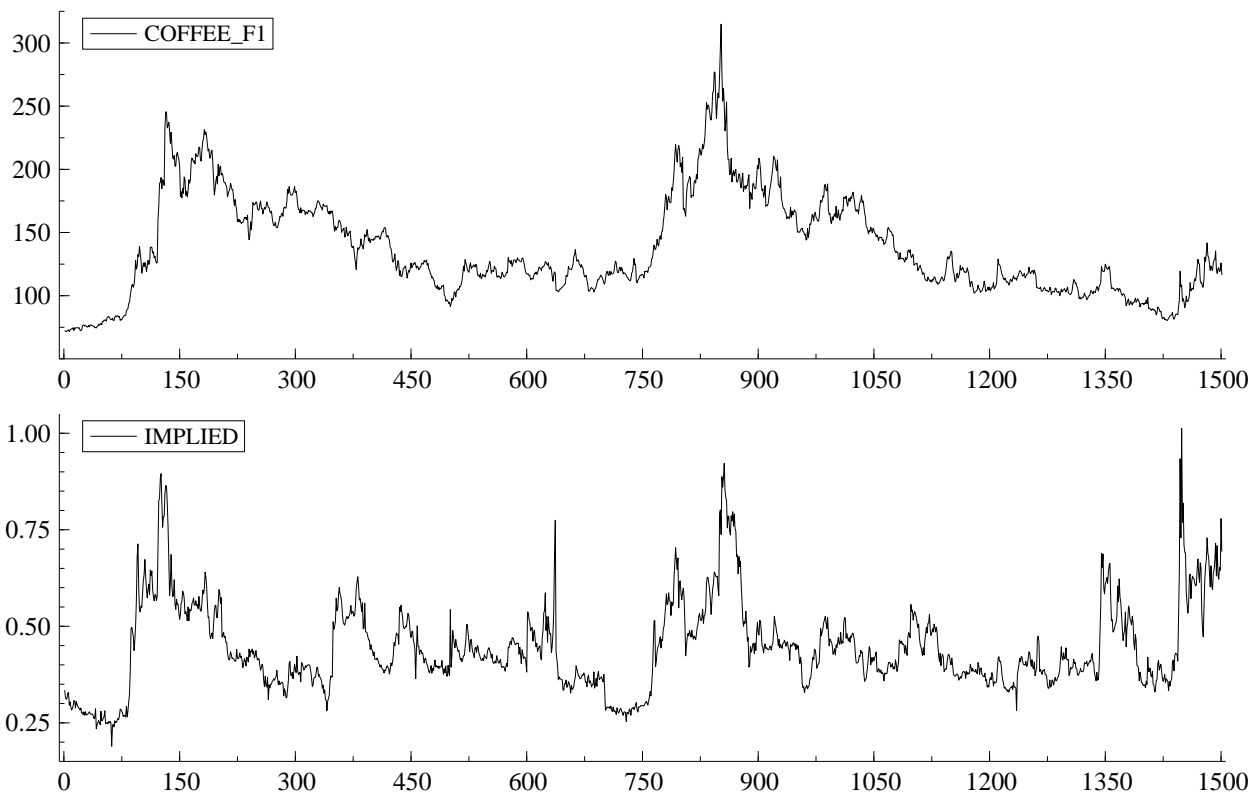


Figure 2: Nearby future coffee daily prices and implied volatility. The time period is 3/1/1994 - 30/12/1999.

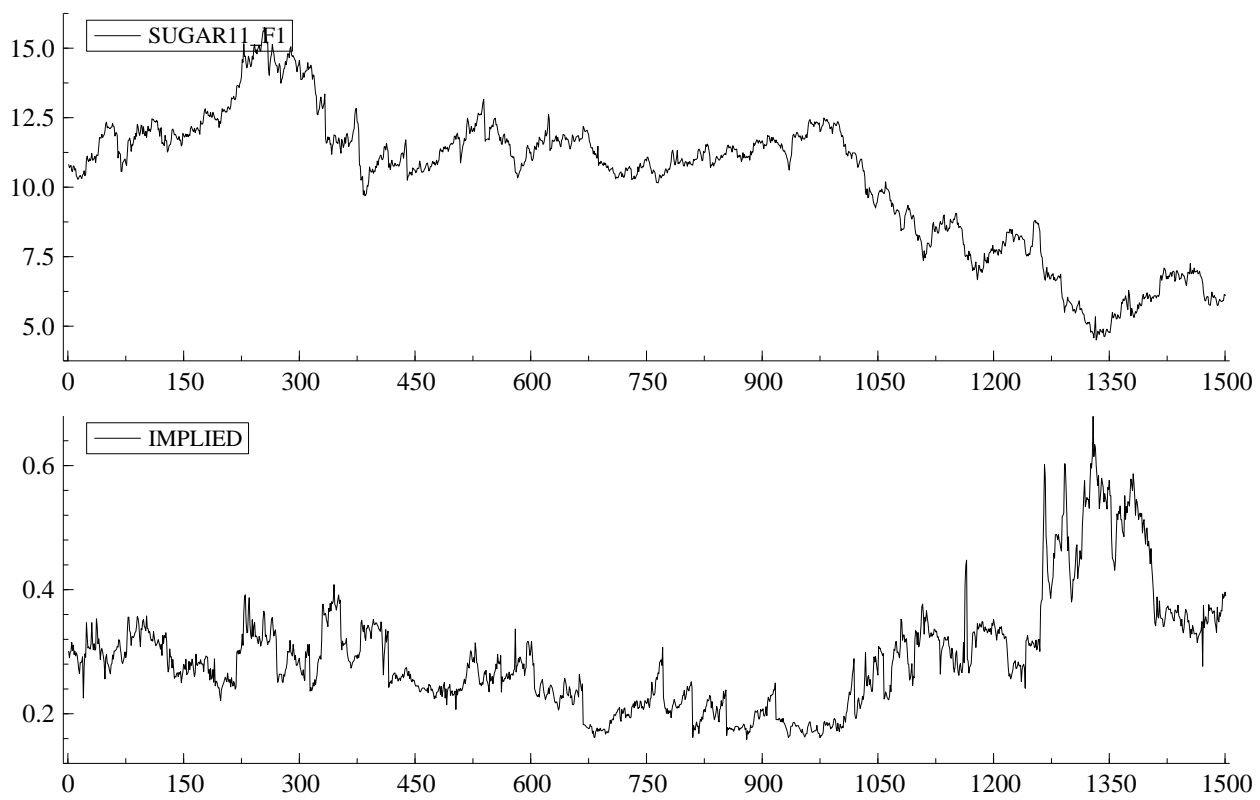


Figure 3: Nearby future sugar daily prices and implied volatility. The time period is 3/1/1994 - 30/12/1999.