Limit Moves as Censored Observations of Equilibrium Futures Price in GARCH Processes

I.G. Morgan and R.G. Trevor

18 July 1997

I.G. Morgan
School of Business
Queen's University
Kingston, Ontario
Canada K7M 2M8

morganig@qucdn.queensu.ca

R.G. Trevor
Centre for Studies in Money, Banking and Finance
Macquarie University
Sydney, New South Wales
Australia 2109

robt@cmbf.mq.edu.au

Abstract

We develop and evaluate an algorithm for estimating GARCH models for time series containing some censored observations. Motivation for the algorithm comes from certain futures market contracts and some equity markets with price limits which constrain the maximum allowable movement in price on a given trading day. When a limit is reached trading stops and the equilibrium price is not observed. The statistical problem is to maximize the likelihood function by replacing the unobservable squared error terms with their expected values. We evaluate the performance of the algorithm by extensive simulation and apply it to Treasury bill futures data from a period of high volatility and frequent limit moves.

Keywords: Price limits, censored observations, rational expectations, GARCH, EM algorithm.

Acknowledgement. We thank Duncan Murdoch and participants in the Queen's University, Macquarie University, Université Laval and University of Ottawa workshops in finance, and the Northern Finance Association Meetings in London, Ontario, for suggestions. The second author gratefully acknowledges support from the Visiting Scholar Program at Queen's University School of Business and from Macquarie Investment Management Limited, Axiom Funds Management and the Australian Research Council under Collaborative Research Grant C595301128. We both acknowledge support from SHRCC.

Limit Moves as Censored Observations of Equilibrium Futures Price in GARCH Processes

Abstract

We develop and evaluate an algorithm for estimating GARCH models for time series containing some censored observations. Motivation for the algorithm comes from certain futures market contracts and some equity markets with price limits which constrain the maximum allowable movement in price on a given trading day. When a limit is reached trading stops and the equilibrium price is not observed. The statistical problem is to maximize the likelihood function by replacing the unobservable squared error terms with their expected values. We evaluate the performance of the algorithm by extensive simulation and apply it to Treasury bill futures data from a period of high volatility and frequent limit moves.

Keywords: Price limits, censored observations, rational expectations, GARCH, EM algorithm.

1. Introduction

Futures exchanges often impose restrictions on the maximum allowable change in price in a given day before trading is halted. Brennan (1986) argued that such limits are partial substitutes for margin requirements and that successful futures contracts are those that survive in a competitive market because they minimize the cost of trading for participants. Price limits, by hiding the equilibrium price, reduce the incentive to renege on contract commitments after a large adverse price movement has occurred. By doing so, they reduce the overall cost of trading. Brennan's model predicts that price limits will be ineffective if traders can use the prices of related financial instruments to infer the equilibrium price of the futures contract in which trading has been halted. This means that futures exchanges will apply limit rules to some commodity contracts and not to others. For those contracts subject to limit rules, traders must estimate the unobservable equilibrium price from the observable censored price and researchers must develop models to take account of the censoring of some of the observations in the time series.

In the analysis of time series of futures prices containing some censored observations, an immediate question a researcher faces is how to deal with the estimation problem the censoring poses. An obvious and common choice is to adopt a naive approach and ignore the problem by treating the observed censored observations as if they were actually the equilibrium prices. We call this the Ostrich algorithm [Davidson and MacKinnon (1993)].

The Ostrich algorithm was one of two approaches McCurdy and Morgan (1987) used for analysis of data from currency futures markets in a period when the Chicago Mercantile Exchange relied on tight price limits for allowable trading ranges within a day. McCurdy and Morgan's second approach was to throw out all observations contaminated by the censoring, approximately twice the number of days on which the limits occurred because the rate of change in price is the variable under study instead of price itself. Kodres (1988, 1993), wanting to avoid the shortcomings of both of these approaches, developed estimation methods for dealing with data contaminated by the censoring induced by price limits. Two problems with her methods are that they rely on selective use of forward price approximations of the futures price (from covered interest arbitrage) and they are numerically complex. They use the information in the forward price only when the equilibrium price is censored. When these forward price approximations are valid and accurate enough to reveal the hidden equilibrium price for a particular futures contract, the Brennan (1986) model predicts that the futures exchange would find no advantage in

keeping the price limits. An example of the level of numerical complexity is that although the Kodres (1993) model is for univariate time series, it requires many numerical evaluations of the bivariate normal distribution function.

We use the same model as McCurdy and Morgan (1987) and Kodres (1993) for conditional heteroscedasticity in prices and present a simpler estimation technique which incorporates the censoring without the additional numerical complexity of the Kodres (1993) method or its requirement for matching forward price data. We develop a rational expectations (RE) algorithm for censored observations in the presence of heteroscedasticity, estimate the model for the rate of change of futures prices and explore the properties of the procedure by simulation.

Our model incorporates GARCH [Engle (1982), Bollerslev (1986)], for the evolution of the heteroscedasticity in asset prices. Our RE algorithm estimates the conditional expectations of the squares of the underlying equilibrium error terms, and simultaneously updates the GARCH process, to reflect the information on the unobserved equilibrium price revealed by the observed censored price. In related work, Trevor (1994) developed an algorithm for estimating the GARCH model when some data are missing as a result of the absence of trading.

We show that the RE algorithm may be interpreted as an EM algorithm, and hence yields maximum likelihood, (ML), estimates under particular assumptions about the data generating process. Alternatively, it may be interpreted as an approximation to a maximum likelihood method when we do not invoke these assumptions. To quantify the accuracy of the approximation we examine the properties of the RE and Ostrich estimation procedures through simulation of time series with some censored data, comparing the estimates with those obtained from the (unobservable) original or underlying process. While Kodres (1993) did not evaluate the advantages of her numerically complex procedure over the Ostrich algorithm, we show that our straightforward RE algorithm gives results that are superior to those of this commonly used method of ignoring the censoring problem. We demonstrate the use of our algorithm to model US Treasury bill futures data from a high volatility period containing frequent limit moves.

Section 2 presents the details of the computations needed for censored futures price data subject to some

form of heteroscedasticity. The particular feature of the censoring is that the rate of change of futures price is not directly observable on the day the price reaches the limit nor on the following day. Section 3 presents a Monte Carlo investigation of the properties of the Ostrich and RE algorithms and Section 4 applies the model to Treasury bill futures from the period in which the Federal Reserve System targeted the money supply in its operations. Section 5 summarizes the paper and discusses possible applications of the work.

2. Estimation of a GARCH Model with Censored Observations

In this section, we present the adjustments necessary when the equilibrium futures price is unobservable either at day *t* or at day *t-1*. In either case the equilibrium rate of change of futures price cannot be computed directly from the data in the usual way. Given the censoring of price, the estimation problem is to infer the expected equilibrium price and corresponding rate of change of price. The algorithm must accommodate the adjustment by replacing the squared error terms in the likelihood function with their corresponding conditional expectations. Our RE algorithm is motivated by the EM algorithm [Dempster, Laird and Rubin (1977)]. Previous research using the EM algorithm for censored data has been with homoscedastic models: Blight (1970) formulates and estimates a similar model for data that may be truncated beyond upper and lower limits, Hasselblad, Stead and Galke (1980) use similar expressions for grouped data, and McGill (1995) extends the analysis to multivariate normal data. In related work, Calzolari and Fiorentini (1997) examine a Tobit model with GARCH errors by Monte Carlo methods.

2.1 Impact of censored observations

In order to understand the impact of censored observations on the estimation of the GARCH models used by McCurdy and Morgan (1987) and Kodres (1993), it is helpful to briefly consider the specification and estimation of such models.

Let

F_t be the equilibrium futures price at time *t*

y_t be the rate of change of the equilibrium futures price from *t-1* to *t*

μ be the expected rate of change of the equilibrium futures price from *t-1* to *t*

σ_t^2 be the variance of y_t conditional on equilibrium price information,

$I_{t-1} = \{F_{t-1}, \ldots F_1, \sigma_1^2\}$, available at t-1.

Then a simple GARCH(1,1) specification for the futures price process is

$$y_t = \frac{F_t}{F_{t-1}} - 1 = \mu + \varepsilon_t, \tag{1}$$

where $\varepsilon_t \left| I_{t-1} \right. \sim N(0, \sigma_t^2)$ and

$$\sigma_{t+1}^2 = \alpha_0 + \alpha_1 \varepsilon_t^2 + \alpha_2 \sigma_t^2. \tag{2}$$

When the equilibrium prices are observable, Maximum Likelihood estimation involves maximizing the log-likelihood function, conditional on σ_1^2,

$$\sum_{t=1}^{T} \ell_t = -\frac{1}{2} \sum_{t=1}^{T} [\ln(2\pi) + \ln(\sigma_t^2) + \frac{\varepsilon_t^2}{\sigma_t^2}]. \tag{3}$$

However, suppose that at time $t = j$ the price limit rules of the futures exchange prevent the equilibrium price, F_j, from being observed. Instead, the limited, or censored, price, F_j^*, is observed. The censoring prevents the computation of both y_j and y_{j+1}. More importantly, it prevents the computation of ε_j^2 and ε_{j+1}^2, and hence σ_i^2 for all $i > j$. The recursive nature of the conditional variance specification of the

GARCH model means that any contamination of today's data will feed through to all future estimates of conditional variance.

One way to address the problem introduced by censored observations is to ignore it, as in McCurdy and Morgan (1987) and use the observed price, F_t^*, instead of the equilibrium price, F_t, in the estimation of the model. Clearly, this Ostrich algorithm will result in estimates of the conditional variance, σ_t^2, which are biased downwards by the truncation of movements in the futures price by a binding price limit rule.

Another approach to this problem is to recognize that the censoring of the futures price provides partial information about the value of the equilibrium price. Kodres (1993) incorporates this information into her estimation algorithm by numerically integrating over the relevant part of the bivariate normal density for each affected observation and each iteration of the maximization algorithm.

Note that the GARCH model likelihood function given by equation (3) can be written completely in terms of squares of the news innovations, ε_t. (Simply substitute equations (1) and (2) into (3).) This suggests an alternative algorithm that replaces the unobservable elements of the likelihood function by their expectations conditional on the information available to traders. To implement the RE algorithm in this context we need to consider expectations of squares of the equilibrium news innovations conditional on the observed innovations.

2.2 The information revealed by limit moves

In this subsection we classify the information revealed by limit moves but leave the task of quantifying the information to a later subsection. We argue that there are only three basic cases. A fourth case, involving observed equilibrium prices at both *t-1* and *t*, completes the classification.

In practice, periods of high volatility of price are likely to result in futures prices that reach price limits on a number of consecutive days. However, close inspection of the likelihood function and GARCH specification, given by equations (1) through (3), shows that the nature of the recursion in the model means that all possible sequences of limit moves can be accounted for by considering what happens on

any two consecutive days. The complete event space is thus spanned by nine possibilities. Table 1 lists the nine possible sequences for censored and uncensored prices at *t-1* and *t*. We discuss selected rows of Table 1 in the following paragraphs.

Table 1: Consecutive limit moves and news

Price limits hit		Observed price		News at time t
t-1	t	t-1	t	
upper	upper	$F^*_{t-1} \le F_{t-1}$	$F^*_t \le F_t$	none
upper	none	$F^*_{t-1} \le F_{t-1}$	F_t	$\varepsilon^*_t \ge \varepsilon_t$
upper	lower	$F^*_{t-1} \le F_{t-1}$	$F^*_t \ge F_t$	$\varepsilon^*_t \ge \varepsilon_t$
none	upper	F_{t-1}	$F^*_t \le F_t$	$\varepsilon^*_t \le \varepsilon_t$
none	none	F_{t-1}	F_t	ε_t
none	lower	F_{t-1}	$F^*_t \ge F_t$	$\varepsilon^*_t \ge \varepsilon_t$
lower	upper	$F^*_{t-1} \ge F_{t-1}$	$F^*_t \le F_t$	$\varepsilon^*_t \le \varepsilon_t$
lower	none	$F^*_{t-1} \ge F_{t-1}$	F_t	$\varepsilon^*_t \le \varepsilon_t$
lower	lower	$F^*_{t-1} \ge F_{t-1}$	$F^*_t \ge F_t$	none

In the case of a single down limit move, it must be that the observed censored price $F^*_t \ge F_t$ the unobserved equilibrium price. Hence the news revealed by the observed censored price is $\varepsilon^*_t \ge \varepsilon_t$ the news that would be revealed by the equilibrium price; while with a single up limit move, we know that $F^*_{t-1} \le F_{t-1}$ and $\varepsilon^*_t \le \varepsilon_t$.

Now consider the case of a lower limit move at time t, preceded by an upper limit move at time *t-1* (i.e. the third row of Table 1). Then we have $F^*_{t-1} \le F_{t-1}$ and $F^*_t \ge F_t$, resulting in $\varepsilon^*_t \ge \varepsilon_t$, by equation (1). Thus

the news revealed by an up-down limit move is analogous to the news revealed by a single down limit move (the sixth row of Table 1). In both cases, the censored news innovation is greater than the news innovation that would be revealed by the equilibrium price.

Next, consider consecutive limit moves in the same direction (the first and last rows of Table 1). In such cases, the censored news innovation provides no information on the news innovation that would be revealed by the equilibrium price. While we know that today's equilibrium price must lie above the censored price when there have been consecutive up limit moves, we cannot deduce whether the observed rate of change in futures price lies above or below the equilibrium rate of change in the futures price. Hence the censored news innovation is uninformative about the equilibrium news innovation.

Further consideration of the censored news innovation, ε_t^*, in Table 1 shows that the eight cases involving censoring reduce to three possibilities. The censored value ε_t^* provides an upper bound for the equilibrium value ε_t if the censoring sequence for two consecutive days is either up-none, up-down, or none-down, (rows 2, 3 and 6); it provides a lower bound if the sequence is either none-up, down-up, or down-none, (rows 4, 7 and 8); and it is uninformative if the sequence is either up-up or down-down.

2.3 The RE algorithm

In the simple GARCH specification, the log-likelihood function (equation (3)) may be written entirely in terms of the squared equilibrium news innovations, ε_t^2. (More general specifications may also involve terms in ε_t, but they involve straightforward extensions of the results we derive below.) When price limits induce censoring of the observable futures prices, we have shown that any sequence of limit moves will result in one of only three possibilities. The censored news innovation will provide either an upper or lower bound on the equilibrium innovation or be uninformative. The RE algorithm involves replacing any unobservable squared equilibrium news innovation, ε_t^2, by its expectation conditional on the information revealed to the trader by the corresponding observed quantity, ε_t^*.

We now distinguish between the previously defined (unobservable) information set I_t, containing the

equilibrium prices and expressible in recursive form as $I_1 = \{F_1, \sigma_1^2\}, I_t = I_{t-1} \cup \{F_t\}$, and the trader's information set A_t, containing the observed prices. This set may be expressed as $A_1 = \{F_1, \sigma_1^2\}$ if F_1 is not censored, $A_1 = \{F_1^*, \sigma_1^2\}$ otherwise; $A_t = A_{t-1} \cup \{F_t\}$ if F_t is not censored, $A_t = A_{t-1} \cup \{F_t^*\}$ otherwise.

Consider the case where the current observed innovation, ε_t^*, provides an upper bound on the equilibrium news innovation, ε_t, that is, when the censoring sequence has been up-none, up-down or none-down. In computing the ℓ_t term of the log-likelihood in equation (3) we replace the unobservable ε_t^2 term by its expectation conditional on the observed, censored innovation, ε_t^*. In other words, we replace ε_t^2 by $E[\varepsilon_t^2 \,|\, A_t]$, where $\varepsilon_t \notin A_t$, but $\varepsilon_t \le \varepsilon_t^* \in A_t$. However, to compute the ℓ_{t+1} term of the log-likelihood we need to compute σ_{t+1}^2, which depends on the unobservable ε_t^2. We again replace ε_t^2 by $E[\varepsilon_t^2 \,|\, A_t]$, where $\varepsilon_t \notin A_t$, but $\varepsilon_t \le \varepsilon_t^* \in A_t$.

Let $z_t = \varepsilon_t / \sigma_t$, $\phi(z)$ the unit normal density function and $\Phi(z)$ the corresponding distribution function. Then[1], when $\varepsilon_t \notin A_t$, but $\varepsilon_t \le \varepsilon_t^* \in A_t$,

$$E[\varepsilon_t^2 \,|\, A_t] = \sigma_t^2 E[z_t^2 \,|\, z_t \le z_t^* = \frac{\varepsilon_t^*}{\sigma_t}] = \sigma_t^2 [1 - z_t^* \frac{\phi(z_t^*)}{\Phi(z_t^*)}] = \sigma_t^2 [1 - \frac{\varepsilon_t^* \phi(\varepsilon_t^* / \sigma_t)}{\sigma_t \Phi(\varepsilon_t^* / \sigma_t)}]. \qquad (4)$$

Similar arguments apply when the current observed innovation, ε_t^*, provides a lower bound on the equilibrium news innovation, ε_t; that is, when the censoring sequence has been down-none, down-up or none-up. In computing σ_{t+1}^2 and the ℓ_t term of the log-likelihood in equation (3) we replace the

---

[1] Note that to replace ε_t^2 by its rational expectation we need $E[z_t^2 \,|\, z_t \le z_t^*]$, as in equation (4), (for example, Raiffa and Schlaifer (1961, page 232)), rather than $Var[z_t^2 \,|\, z_t \le z_t^*]$, as in Maddala (1983, pages 365-367).

unobservable ε_t^2 term by its expectation conditional on the censored innovation, ε_t^*. In other words, when $\varepsilon_t \notin A_t$, but $\varepsilon_t \geq \varepsilon_t^* \in A_t$ we replace ε_t^2 by

$$E[\varepsilon_t^2 \Big| A_t] = \sigma_t^2 E[z_t^2 \Big| z_t \geq z_t^* = \frac{\varepsilon_t^*}{\sigma_t}] = \sigma_t^2 [1 + z_t^* \frac{\phi(z_t^*)}{1 - \Phi(z_t^*)}] = \sigma_t^2 [1 + \frac{\varepsilon_t^* \phi(\varepsilon_t^* / \sigma_t)}{\sigma_t \Phi(-\varepsilon_t^* / \sigma_t)}]. \quad (5)$$

Clearly, when no limit move has occurred at either time *t* or *t-1*, the equilibrium news innovation $\varepsilon_t \in A_t$ is observed and the computation of the ℓ_t term of the log-likelihood function proceeds in the usual manner.

This leaves the case of consecutive limit moves in the same direction; up-up or down-down. Then the censored news innovation, ε_t^*, is uninformative about the equilibrium news innovation, ε_t. This corresponds to GARCH with a missing observation considered by Trevor (1994), and we apply the algorithm developed there: since the current observation is uninformative, we exclude the corresponding ℓ_t term from the log-likelihood function. To compute the ℓ_{t+1} term of the log-likelihood function we need σ_{t+1}^2, which depends on the unobservable ε_t^2. We replace this by its conditional expectation $E[\varepsilon_t^2 \Big| A_t]$, where $\varepsilon_t \notin A_t, \varepsilon_t^* \in A_t$, but ε_t^* uninformative, so

$$E[\varepsilon_t^2 \Big| A_t] = E[\varepsilon_t^2 \Big| A_{t-1}] = \sigma_t^2, \quad (6)$$

since the censored observation reveals no information that could be used to update the conditional variance.

The expectations in equations (4), (5) and (6) are the only additional items required to implement our proposed RE algorithm for estimating a GARCH model in the presence of censored observations. None of these expressions involve a significant computational burden, since simple, but highly accurate,

polynomial approximations are readily available for the normal distribution function.[2]

While the above derivations are for the simple GARCH specification where the conditional variance equation does not involve expressions in the level of the news innovation, and the conditional distribution is normal, similar results may be derived for other specifications. Extensions to specifications in which the conditional variance is also a function of the level of the news innovation, and/or the conditional distribution is Student-t, are particularly straightforward. Such extensions also result in expressions which are computationally simple.

Our proposed RE algorithm is:

- initialization: supply initial values for the parameters of the model given by equations (1) and (2), i.e. the vector $(\mu,\alpha_0,\alpha_1,\alpha_2)'$ and σ_1^2 to initialize equation (2);

- expectation step: conditional on the current values of the parameter vector, compute the conditional expectations of ε_t^2 (from equations (4), (5) or (6)) for any observations affected by a limit move in the current or previous time period;

- maximization step: conditional on these computed expectations of ε_t^2 for observations affected by censoring, maximize the log-likelihood function of the model (equation (3)) to obtain new estimates of the parameter vector;

- repeat this sequence of expectation and maximization steps until the maximum of the log-likelihood function for the current sequence differs from that of the previous sequence by less than a tolerance of 0.001.

While we choose to develop this algorithm as a repetitive cycle of two steps, we could, in principle, collapse the expectation step into the maximization step. We have found that the two step routine converges well.

---

[2] We use Algorithm AS66 of Applied Statistics(1973).

2.4 Interpretation of the RE Algorithm

There are two ways of interpreting this algorithm. The first is as an EM algorithm. The unobservable equilibrium price probability density function (pdf) for the specification given in equations (1) and (2) is

$$f(y_T,\cdots,y_1;\theta\left|\sigma_1\right.)=\prod_{t=1}^{T} f_t(y_t;\theta\left|y_{t-1},\cdots,y_1,\sigma_1\right.)=\prod_{t=1}^{T} f_t(y_t;\theta\left|I_{t-1}\right.), \tag{7}$$

where $\theta'=(\mu,\alpha_0,\alpha_1,\alpha_2)'$. That is, the joint pdf is constructed from the product of the conditional pdfs.

Since each of the conditional pdfs is normal (and therefore a member of the exponential family), each may be expressed in terms of its sufficient statistics. To simplify the exposition, consider the case where $\mu=0$, so that $\varepsilon_t=y_t$ and θ' is reduced to $(\alpha_0,\alpha_1,\alpha_2)'$. Then

$$f_t(y_t;\theta\left|I_{t-1}\right.)=b(y_t\left|I_{t-1}\right.)\exp\left[\Omega_t(\theta\left|I_{t-1}\right.)\cdot s(y_t\left|I_{t-1}\right.)\right]\Big/a(\Omega_t(\theta\left|I_{t-1}\right.)), \tag{8}$$

where $b(y_t\left|I_{t-1}\right.)=1$

$$\Omega_t(\theta\left|I_{t-1}\right.)=1\Big/\sigma_t^2(\theta\left|I_{t-1}\right.)$$

$$s(y_t\left|I_{t-1}\right.)=-y_t^2/2=-\varepsilon_t^2/2$$

$$a(\Omega_t(\theta\left|I_{t-1}\right.))=\sqrt{2\pi}\,\sigma_t(\theta\left|I_{t-1}\right.).$$

Dempster, Laird and Rubin (1977) show that for the exponential family, repeated cycling from iteration $p, p+1,$ …, in the following EM algorithm produces ML estimates:

E-step compute $s^{(p)}(y_t\left|I_{t-1}\right.)=E[s(y_t\left|I_{t-1}\right.)\left|A_T;\Omega_t(\theta^{(p)}\right|I_{t-1}\right.)]$

M-step compute $\Omega_t(\theta^{(p+1)}\left|I_{t-1}\right.)$ by maximizing

$$\sum_{t=1}^{T} -\log(a(\Omega_t(\theta|I_{t-1}))) + \Omega_t(\theta|I_{t-1}) \cdot s^{(p)}(y_t|I_{t-1}) .$$

It is a matter of computational convenience as to whether these steps are computed recursively or simultaneously.

In the presence of censoring, our RE algorithm for estimating equations (1) and (2) is not the same as this EM algorithm when the equation for the conditional variance (equation (2)) depends on past values of the unobservable *equilibrium* futures price (or innovation). In such situations $\Omega_t(\theta|I_{t-1})$ depends on the unobservable information I_{t-1} rather than the observable A_{t-1}, so that the expectations required by the E-step do not correspond to the expectations given by equations (4), (5) and (6). If instead we assume that the imposition of price limits by an exchange alters the data generating process for equilibrium prices so that the conditional variance process depends only on information observable to traders, we replace equation (2) by

$$\sigma_{t+1}^2 = \alpha_0 + \alpha_1 E(\varepsilon_t^2|A_t) + \alpha_2 \sigma_t^2 . \qquad (9)$$

Then $\Omega_t(\theta|\bullet)$ becomes $\Omega_t(\theta|A_{t-1})$ and the required E-step expectations collapse to those given by equations (4), (5) and (6). Our RE algorithm would then be an EM algorithm in the sense of Dempster et al (1977) and would produce ML estimates of equations (1) and (9).

Since $\varepsilon_t^2 = E(\varepsilon_t^2|I_t)$ (in equation (2)) and $E(\varepsilon_t^2|A_t)$ (in equation (9)) differ only when F_t and /or F_{t-1} are unobservable (censored), we can also think of our RE algorithm as producing approximate ML estimates for equations (1) and (2). That is, the RE algorithm produces ML estimates of equations (1) and (9) and approximate ML estimates of equations (1) and (2).[3] The accuracy of approximation will depend on the degree to which the conditional variance process in equation (9) approximates that in equation (2). This is likely to depend on both the severity of censoring and the parameters of the model, and may be evaluated by means of a Monte Carlo analysis.

---

[3] Calzolari and Fiorentini (1997) also discuss the use of approximate ML estimates.

3. Monte Carlo analysis for algorithm comparisons

In this section we present the experimental design for simulating and estimating a simple univariate GARCH model of random rate of change of price data subject to occasional censoring. In Section 2.4 we explained how application of the RE algorithm to data generated by equations (1) and (9) is an application of the EM algorithm and hence will produce ML estimates. Suppose, instead, that the underlying, unobservable, equilibrium process actually corresponds to equations (1) and (2), so that the imposition of price limits does not distort the conditional variance of the equilibrium price process. Then the RE algorithm approximates this conditional variance by using the observed data. A Monte Carlo analysis in which we generate the equilibrium price process by equations (1) and (2) and then censor some of the data, will allow us to examine how accurately the RE algorithm tracks the true conditional variance process. The Monte Carlo analysis also allows us to compare the RE algorithm with the commonly used Ostrich algorithm in this respect.

3.1 Research design

Futures and some equity exchanges, notably the Tokyo Stock Exchange, impose price limits in a variety of ways. Some are intradaily and some are daily, some are fixed and some are escalating so that the limits are gradually relaxed.[4] Our RE algorithm applies to all of these but we need to choose one scheme for the Monte Carlo study. We choose to mimic a daily observation interval with fixed price limits for two reasons. First, we wish to subject our RE algorithm to conditions that are stringent in the sense that we do not relax the limits after a limit move occurs. Second, even if an Exchange sets price limits that are rarely effective under normal circumstances, the clustering typical of financial data will generate enough volatility for the price to reach the limit. For example, Phillips and Weiner (1996) report that the New York Mercantile Exchange (NYMEX) crude oil futures price for at least one contract hit the limit on 29 days in the seven months of the Persian Gulf crisis starting in August 1989. In 1997 disruption of the supply of palladium has led to several days of limit moves on the NYMEX. Palladium futures provide a good illustration of clustering of high volatility; of all limit moves from 1980 through June 1997, approximately two thirds occurred in 1980; and 1997 is the year with the second highest frequency of

[4] The Chicago Board of Trade expands the limits on many of its financial futures to 1.5 times the usual limit on the day following a limit move. We could easily simulate this process, but it is only one of many possible institutional choices. The New York Mercantile Exchange, for example, often uses a factor of 2.0.

limit moves. In Section 4 we apply the algorithm to data from Treasury bill futures in a period in which the Chicago Mercantile Exchange imposed fixed daily limits and these limits affected trading because volatility was high.

In the system of equations (1) and (2) we choose parameter values that are typical of daily futures rates of change of price. Risk premiums in daily futures data are generally small (for example, McCurdy and Morgan (1987)), and unconditional average daily rates of change in price are generally insignificantly different from zero. Therefore, we set $\mu = 0$ in equation (1). Lumsdaine (1995) shows that the choice of μ does not affect the sampling distribution of the estimates of the other parameters. We choose four sets of these conditional variance parameters to be representative of daily data and to reflect the levels of persistence $(\alpha_1 + \alpha_2)$ of shocks to conditional variance common in daily data. We choose α_0 so that the unconditional variance is unity.[5] Table 2 shows the parameter sets.

For each parameter set A, B, C and D we apply four different thresholds of censoring, corresponding to *a = 1.5, 2.0, 2.5*, and *3.0* unconditional standard deviations. The choice *a = 2.0* is a close approximation to the censoring threshold in the Treasury bill futures application in Section 4. The larger values, *a=3.0* and *a=2.5,* correspond to low frequency of limited prices more typical of many other futures contracts. To subject the RE algorithm to more stringent conditions in case its approximation to the true process deteriorates with limit move frequency, we include the relatively heavy censoring level *a = 1.5*.

---

[5] Lumsdaine (1995) shows that the choice of α_0 does not affect the estimates of α_1 or α_2.

Table 2: Input parameter values

parameter set	A	B	C	D
μ	0	0	0	0
α_0	0.05	0.01	0.01	0.01
α_1	0.05	0.04	0.14	0.09
α_2	0.90	0.95	0.85	0.90

For each of the 16 sets of conditions (four parameter sets and four censoring thresholds), simulate 1100 uncensored observations of the equilibrium GARCH process (equations (1) and (2)) from 1100 unit normal random numbers, using the same initial seed for all sets. Assume the futures exchange implements a fixed limit move rule with no widening of the limits after one or more limit moves have occurred. Generate the censored data time series by copying the uncensored time series but replacing observations $y_t \geq a$ with the value *a*, and observations $y_t \leq -a$ with the value *-a*. In estimation, discard the first 100 observations before fitting the GARCH model. Flag the censored observations distinguishing between positive and negative limit moves with an indicator variable $L_t = +1, -1$, respectively. To complete the censored data time series, when a sequence of limit moves starts at *t* and ends with an uncensored price at *t+k*, obtain the observations at time *t* in the censored series, replace the uncensored y_{t+k} with the censored data y^*_{t+k} from the equilibrium price series, as

$$y^*_{t+k} = 100 \times \left[\frac{\prod_{j=0}^{k} (1 + y_{t+j} / 100)}{\prod_{j=0}^{k-1} (1 + aL_t / 100)} - 1 \right]. \qquad (10)$$

For example, if *k=1*, following an upper limit move at *t*,

$$y^{*}_{t+1} = 100\times\left[\frac{(1+ y_t/100)(1+ y_{t+1}/100)}{(1+ a/100)} - 1\right]. \qquad (11)$$

For each replication of the set of underlying data and censored data estimate the parameters $(\mu, \alpha_0, \alpha_1, \alpha_2)$ of the GARCH model[6] by three algorithms:

1) by the RE algorithm to the censored data;
2) by the Ostrich algorithm (i.e., maximum likelihood to the censored data, ignoring the presence of censoring);
3) by maximum likelihood to the ordinarily unobservable equilibrium data. This is, of course, only possible in a Monte Carlo experiment.

Replicate each set 1000 times. If any one of these three numerical estimations fails to converge, abandon all three, store the information about which algorithm or algorithms failed, treat the set as unsuccessful and replace it. The decision to replace unsuccessful replications follows Engle and Ng (1993).

3.2 Criteria for evaluation of results

In Section 2.4 we emphasized that the crucial question for this Monte Carlo study is whether the RE algorithm produces conditional variance estimates that closely approximate the conditional variance of the unobservable equilibrium data given by equation (2). To evaluate the performance of these algorithms we use three main criteria. First, we examine the accuracy of each algorithm in estimating the time series of the standardized error terms in (1). This involves the simple linear regression,

$$\varepsilon_t/\sigma_t = b_0 + b_1\,\varepsilon_t/\hat{\sigma}_{it} + u_{it} \qquad (12)$$

[6] In all estimations for this paper we set the initial variance equal to the unconditional variance implied by the parameter estimates from the previous iteration of the non-linear optimization routine.

for algorithm *i, i = 1,2,3*. The dependent variable equation (12) is the actual standardized residual from the simulation. If the estimated conditional standard deviation for a given algorithm, $\hat{\sigma}_{it}$, represents the actual conditional standard deviation accurately, b_0 should be zero and b_1 should be *1*. These parameters will be consistently estimated even if the standardized residuals are autocorrelated. Under the strict null hypothesis these standardized residuals are independently and identically distributed as unit normal. Consequently, we stack the residuals from each replication into a single regression analysis.[7]

Secondly, we modify a sign test from Amin and Ng (1994) to evaluate the comparative accuracy of pairs of algorithms in estimating conditional variance. Since GARCH processes are recursive, errors of estimation of conditional variance may not be independent from one observation to the next, but cumulative errors of estimation from each entire replication must be mutually independent.

Define, for replication k, the cumulative difference in absolute estimation error, d_{ij}^k, and an indicator variable $I(d_{ij}^k)$ that takes the value *1* if $d_{ij}^k < 0$ and zero otherwise. Then r is the frequency, in $n = 1000$ replications, with which algorithm i (the first of a pair of algorithms) is more accurate in estimating the true conditional variance than algorithm j (the second of the pair), as in

$$\begin{aligned} d_{ij}^k &= \sum_{t=1}^{T} \left|\hat{\sigma}_{it}^2 - \sigma_t^2\right| - \left|\hat{\sigma}_{jt}^2 - \sigma_t^2\right|, \\ r &= \sum_{k=1}^{n} I(d_{ij}^k), \\ Z &= (r - n/2)/(n/4)^{1/2}. \end{aligned} \tag{13}$$

The test statistic Z assumes the null hypothesis of equal accuracy and is asymptotically unit normal.

Our third criterion is the difference in log-likelihoods for each algorithm. We use mean paired differences to summarize the information from all replications. We focus on the differences between the ML and the other two algorithms to detect which of these two is closest to the ML likelihood.

---

[7] Each of the 1000 replications generates 999 $\varepsilon_t/\hat{\sigma}_{it}$ values for a given algorithm, since the estimates are conditioned on the initial value σ_1. This yields a total of 999,000 observations for each algorithm.

3.3 Simulation results

Table 3 describes the characteristics of the simulated data samples for the four sets of parameter values. The rows in each panel are in order from lightest to heaviest censoring levels. To save space, we show the sample distribution statistics for the raw data corresponding to the censoring threshold *3.0* only. Minor differences in the raw data for the other three thresholds are due to the slightly different sets of *1000* successful replications. There are two main features in the sample variance data. First, as the level of censoring becomes more severe the mean sample variance of the censored data falls, as must happen. Second, for parameter set C the raw data mean sample variance is considerably lower than the unconditional variance of *1.0*.

One replication illustrates some of the small sample problems of GARCH process simulation. The maximum sample variance in parameter sets C (*24.9*) and D (*8.7*) are for the same sets of random numbers and are several times the input unconditional variance of *1.0*. In other sets of *1000* replications, both with the same random number generator and with a different generator, similar outliers occurred. It is clear that huge sample sizes per replication would be necessary to reduce or avoid the effects of occasional samples such as these.

Table 3 : Description of simulated data

	censoring threshold	minimum	Q1	median	Q3	maximum	mean
			Parameter set A				
sample variance	none	0.768	0.932	0.992	1.060	1.386	0.999
	3.0	0.768	0.927	0.985	1.048	1.321	0.991
	2.5	0.767	0.913	0.966	1.024	1.284	0.972
	2.0	0.746	0.868	0.912	0.962	1.178	0.916
	1.5	0.673	0.754	0.786	0.823	1.005	0.789
limit move frequency	3.0	0	2	3	5	17	4
	2.5	3	10	14	18	43	14
	2.0	21	40	48	56	100	49
	1.5	93	129	143	157	220	144
			Parameter set B				
sample variance	none	0.561	0.834	0.957	1.119	2.373	0.993
	3.0	0.561	0.830	0.945	1.099	2.009	0.977
	2.5	0.561	0.817	0.926	1.067	1.874	0.951
	2.0	0.559	0.784	0.873	0.993	1.638	0.893
	1.5	0.522	0.691	0.756	0.843	1.467	0.770
limit move frequency	3.0	0	2	4	8	66	6
	2.5	0	8	14	23	117	17
	2.0	5	32	45	63	198	50
	1.5	51	107	133	167	331	139
			Parameter set C				
sample variance	none	0.288	0.527	0.684	1.017	24.913	0.926
	3.0	0.288	0.523	0.659	0.926	15.951	0.802
	2.5	0.288	0.513	0.640	0.875	16.044	0.775
	2.0	0.287	0.493	0.600	0.806	11.410	0.716
	1.5	0.282	0.451	0.539	0.711	11.887	0.646
limit move frequency	3.0	0	1	5	18	269	14
	2.5	0	5	13	34	383	24
	2.0	0	16	32	63	437	45
	1.5	9	50	77	124	495	93
			Parameter set D				
sample variance	none	0.380	0.667	0.830	1.124	8.737	0.980
	3.0	0.380	0.662	0.813	1.051	4.409	0.903
	2.5	0.380	0.652	0.790	1.001	4.014	0.865
	2.0	0.379	0.628	0.744	0.923	4.403	0.806
	1.5	0.368	0.562	0.652	0.798	5.068	0.710
limit move frequency	3.0	0	2	5	15	280	12
	2.5	0	7	15	33	332	24
	2.0	1	23	39	70	399	52
	1.5	20	73	104	153	478	118

Note the data in rows marked *none* correspond to the *1000* successful replications for censoring threshold *3.0.* The sample distributions for the slightly different sets of successful replications for the other three thresholds had identical extrema and similar, but not identical, quartiles and means.

Table 4 shows that the failure rates for the RE and Ostrich algorithms are similar, except for parameter set A. Table 5 shows that for the vast majority of cases, the RE algorithm converges in less than ten steps.

3.4 Evaluation of simulation estimates

Table 6 compares the mean paired differences between log-likelihood values across the algorithms. The main focus is on how closely the RE and Ostrich algorithms approximate the ML log-likelihood. However, the first comparison is for the true or actual log-likelihood for the simulated equilibrium data with the estimated value for the full information maximum likelihood algorithm using the complete set of observations for each replication. This comparison is unaffected by censoring. The small differences across the columns in the first panel of Table 6 are due to slightly different sets of failed replications. Direct evidence that full information maximum likelihood estimation of a GARCH model tends to overfit the data is the negative mean difference in all parameter sets, the estimated log-likelihood being greater than the true log-likelihood.

Direct comparison of censored and uncensored data is not appropriate, since the replacement of extreme values in the uncensored data with censored values obviously increases the log-likelihood for those observations. When we compare pairs of algorithms we do so for the subset of observations that are not directly affected by censoring. In the remaining panels of Table 6, we make three explicit comparisons: the ML with each of the true, RE and Ostrich algorithms. Other comparisons may be made by addition or subtraction of appropriate pairs of means in the table.

The third panel of Table 6 shows that the differences between the log-likelihoods of the ML and RE algorithms increase as the level of censoring increases. Except for the most severe level of censoring (*a=1.5)*, these differences are not large relative to the magnitude of the differences between the true and the ML values.

Table 4 : Number of unsuccessful replications, by algorithm

Parameter set	censoring threshold	RE	ML	Ostrich	Total
A	3.0	9	0	1	10
	2.5	10	0	0	10
	2.0	11	0	0	11
	1.5	3	0	1	3
B	3.0	13	9	9	14
	2.5	13	9	11	17
	2.0	8	9	10	13
	1.5	15	9	13	20
C	3.0	0	0	1	1
	2.5	0	0	0	0
	2.0	1	0	1	1
	1.5	4	0	4	5
D	3.0	0	1	1	1
	2.5	2	1	3	3
	2.0	3	1	3	4
	1.5	7	1	5	8

Table 5 : Number of E Steps

Parameter Set	censoring threshold	minimum	Q1	median	Q3	maximum	mean
A	3.0	1	2	2	2	4	2
	2.5	2	2	3	3	6	3
	2.0	2	4	4	5	11	4
	1.5	4	7	8	9	35	8
B	3.0	1	2	2	2	4	2
	2.5	1	2	3	3	6	3
	2.0	1	3	4	4	17	4
	1.5	1	5	6	8	36	7
C	3.0	1	2	3	3	23	3
	2.5	1	2	3	4	32	4
	2.0	1	3	4	5	28	5
	1.5	1	5	6	7	61	7
D	3.0	1	2	2	3	11	3
	2.5	1	2	3	4	13	3
	2.0	1	3	4	5	36	4
	1.5	3	5	6	7	55	7

Table 6 : Simulation mean log likelihood paired differences

			censoring threshold			
data	comparison	parameter set	3.0	2.5	2.0	1.5
complete observations	True - ML	A	-3.42	-3.42	-3.42	-3.42
		B	-3.31	-3.30	-3.31	-3.31
		C	-2.77	-2.77	-2.77	-2.78
		D	-3.03	-3.03	-3.03	-3.04
uncensored subset	True - ML	A	-2.98	-2.60	-2.25	-2.22
		B	-2.89	-2.61	-2.37	-2.16
		C	-2.75	-2.83	-3.02	-3.23
		D	-2.81	-2.84	-2.88	-2.84
	ML - RE	A	0.12	0.63	3.09	12.04
		B	0.27	0.92	3.45	12.52
		C	1.78	3.20	5.83	12.05
		D	1.22	2.52	5.37	13.17
	Ostrich - ML	A	0.23	1.28	7.06	28.83
		B	0.57	2.08	7.99	27.09
		C	1.37	2.40	4.75	9.73
		D	1.60	3.28	7.32	17.48

Note: True refers to the simulated equilibrium data with the true parameter values, as opposed to ML estimates for the same data. ML refers to full information maximum likelihood. For the uncensored subset statistics, other paired differences may be computed by addition or subtraction of appropriate pairs of differences.

In the final panel of Table 6, we compare the Ostrich and ML estimates. With the exception of parameter set C, the Ostrich algorithm produces log-likelihood values which are further away from the ML values than those produced by the RE algorithm (compare panels 3 and 4).

We now turn to the analysis of the accuracy of the estimates of the conditional variances produced by these algorithms. Table 7 shows estimates of the slope coefficients from regression equation (12). The dependent variable is the simulated equilibrium true error term divided by the true conditional standard deviation. The explanatory variable is the true error term divided by the estimate of the conditional standard deviation from each of the three algorithms. Since the total number of observations in each regression in Table 7 is almost one million and interpretation of significance with conventional *p*-values will not be appropriate, we will focus on the coefficient estimates without reference to their standard errors.

In Table 7, if the sets of 1000 successful replications did not vary slightly as a result of different replications failing as the level of censoring changes, the results for the full information maximum likelihood would always be the same for a given parameter set. Lack of independence across censoring levels precludes formal comparisons of performance. Informally, Table 7 reveals that the RE algorithm performs about as well as the full information maximum likelihood estimates. For example, the RE slope coefficient is closer to 1.0 in *9* analyses, the full information slope is closer in *7*. In Section 2.4 we showed that the approximation error introduced by the RE algorithm was due to the differences between equations (2) and (9). These results suggest that those differences are small.

The Ostrich algorithm performs badly in these regressions. For example, its slope coefficient is always smaller than 1.0 and always further from 1.0 than are the slopes of the other two algorithms. Moreover, its performance deteriorates as the severity of censoring increases. A slope less than unity corresponds to underestimating the conditional variance and the underestimation with the Ostrich algorithm is clearly due to its use of the censored data without adjustment for the censoring.

While Table 7 shows the regression estimates of the individual algorithms, Table 8 presents sign tests of the relative accuracy of pairs of algorithms in estimating the true conditional variance. In Table 8, a positive test statistic *Z* means the first algorithm of the pair is more accurate. It is no surprise that

maximum likelihood with full information of the normally unobservable equilibrium data is almost always more accurate than either the RE or the Ostrich algorithm. When only censored data are available, the RE algorithm is clearly more accurate than the Ostrich algorithm: the test statistic is significantly positive at the 5 percent level in *14* analyses and significantly negative in only one analysis.

In summary, the Monte Carlo analysis shows that the RE algorithm performs better than the Ostrich algorithm. In Table 6 it tracked the ML log-likelihood values more accurately, in Table 7 its conditional standard deviation tracked the true simulated data more closely, and the sign test in Table 8 showed that it estimated the conditional variance more accurately.

What is perhaps more important is that the Monte Carlo analysis shows that the RE algorithm produces estimates of the conditional variance that are close to the true equilibrium values given in equation (2). Therefore, if the imposition of price limits does not affect the equilibrium price process, the approximation errors introduced by the RE algorithm will be small. Naturally if, instead, equation (9) holds, then the RE algorithm is an EM algorithm and will produce maximum likelihood estimates.

Table 7 : Evaluating algorithm performance through estimated conditional standard deviation in standardized error terms

$$\varepsilon_t / \sigma_t = b_0 + b_1 \varepsilon_t / \hat{\sigma}_{it} + u_{it}$$

		slope		
censoring levelthreshold	parameter set	RE	ML	Ostrich
3.0	A	0.9993	0.9985	0.9952
	B	0.9993	0.9981	0.9934
	C	0.9968	0.9937	0.9807
	D	0.9987	0.9959	0.9853
2.5	A	1.0013	0.9986	0.9865
	B	1.0014	0.9981	0.9839
	C	0.9989	0.9937	0.9700
	D	1.0014	0.9959	0.9746
2.0	A	1.0085	0.9986	0.9591
	B	1.0079	0.9981	0.9566
	C	1.0026	0.9937	0.9462
	D	1.0066	0.9959	0.9497
1.5	A	1.0331	0.9986	0.8900
	B	1.0296	0.9981	0.8903
	C	1.0125	0.9937	0.8981
	D	1.0223	0.9958	0.8954

Note: The estimates for the intercept range from -0.0002 to 0.0001. The largest standard error is 0.0002, for both the intercept and slope estimates respectively. Each regression involves *999,000* observations.

Table 8 : Sign tests of algorithm comparative accuracy in estimating conditional variance

$$d_{ij}^{k} = \sum_{t=1}^{T} \left|\hat{\sigma}_{it}^{2} - \sigma_{t}^{2}\right| - \left|\hat{\sigma}_{jt}^{2} - \sigma_{t}^{2}\right|,$$

$$r = \sum_{k=1}^{n} I(d_{ij}^{k}),$$

$$Z = (r - n/2)/(n/4)^{1/2}.$$

censoring threshold	parameter set	models compared: ML v RE		ML v Ostrich		RE v Ostrich	
		r	Z	r	Z	r	Z
3.0	A	623	7.78	497	-0.19	435	-4.11
	B	651	9.55	652	9.61	532	2.02
	C	801	19.04	877	23.84	587	5.50
	D	767	16.89	853	22.33	600	6.33
2.5	A	735	14.86	646	9.23	472	-1.77
	B	743	15.37	839	21.44	626	7.97
	C	850	22.14	956	28.84	635	8.54
	D	824	20.49	951	28.52	692	12.14
2.0	A	865	23.09	895	24.92	549	3.10
	B	837	21.31	987	30.80	751	15.88
	C	899	25.24	993	31.18	669	10.69
	D	872	23.53	998	31.50	752	15.94
1.5	A	979	30.30	1000	31.62	558	3.67
	B	954	28.71	1000	31.62	707	13.09
	C	958	28.97	1000	31.62	643	9.04
	D	961	29.16	1000	31.62	716	13.66

Notes: $I(d_{ij}^{k})$ is an indicator variable with the value *1* if $d_{ij}^{k} < 0$ and zero otherwise in replication *k*. *r* is the frequency, in $n = 1000$ replications, with which algorithm *i* (the first of a pair of algorithms) was more accurate in estimating the true conditional variance than algorithm *j* (the second of the pair). The test statistic *Z* assumes the null hypothesis of equal accuracy and is asymptotically unit normal.

4. Application to Treasury bill futures

4.1 Background

From October 1979 to October 1982 the Federal Reserve system targeted the money supply in its operating procedures. Interest rates were exceptionally variable over most of this period. In the Treasury bill futures market prices moved to their limit frequently in spite of an increase, effective 19 June 1980, from *0.50* to *0.60* as the maximum allowable daily change in the I.M.M. index. There were limit moves in the first deferred contract on *57* days out of *778* in the period. In contrast, no limit move has occurred since October 1982 in this market, and the Chicago Mercantile Exchange no longer imposes price limits on the Treasury bill contract.

4.2 Data

We computed rates of change of futures prices for the set of first deferred contracts for the period October 1979 to October 1982, inclusive. Of the *57* limit moves in this period, *9* occurred on days immediately following days with limit moves in the same direction. We define indicator variables for each day of the week and for the first trading day following the H6 announcement of the Federal Reserve Board. These announcements came after the market close, on Thursdays before 8 February 1980 and on Fridays for the rest of the sample period.

4.3 Model specification, estimation and tests

Since conditional variance may vary over the week in a regular pattern, and may depend on the H6 announcement, we define a set of indicator variables D_{jt}, one for each trading day of the week, ($j = 1,2,\ldots 5$) and one, (*j=6*), for the day following the H6 announcement. We adopt a specification due to Baillie and Bollerslev (1989) and choose a GARCH model with indicator variables. A series of specification searches showed that the Monday, Thursday and Friday indicator variables (D_1, D_4, D_5) did not significantly contribute to the explanatory power of the model, as will be shown in Table 10. This resulted in the specification

$$
\begin{aligned}
y_t &= \mu + \varepsilon_t \\
\varepsilon_t | I_{t-1} &\sim N(0, \sigma_t^2) \\
\sigma_{t+1}^2 &= \alpha_0 + \alpha_1 \varepsilon_t^2 + \alpha_2 \sigma_t^2 + \sum_{j=2,3,6} \gamma_j [D_{jt+1} - (\alpha_1 + \alpha_2) D_{jt}].
\end{aligned}
\tag{14}
$$

Table 9 shows the estimates of this model produced by both the RE and Ostrich algorithms. Notice that, ignoring the indicator variables, the parameter estimates are close to those of parameter set B in the Monte Carlo analysis. Moreover, the censoring threshold corresponds to the case *a=2.0* in that analysis. The results in Tables 6, 7 and 8 for parameter set B suggest that the RE algorithm is producing approximate ML estimates with small approximation error if the equilibrium conditional variance process is assumed to be unaffected by the imposition of the price limit rules.

Table 10 provides evidence that the specification of equation (14) is satisfactory since the model and its standardized residuals survive various diagnostic tests including robust conditional moment tests for omitted variables [Wooldridge (1990)]. In particular, the addition of a moving average error term, or day-of-the-week variables, to the conditional mean would not improve the specification. Nor would the addition of the remaining indicator variables improve the conditional variance.

Returning to Table 9, notice that all three of the γ_j parameter estimates are greater in absolute value and significance for the RE algorithm than for the Ostrich algorithm. The coefficient estimate for the H6 announcements, γ_6 ,is particularly strongly affected by the choice of algorithm and the effect of these announcements on Treasury bill futures prices becomes much clearer with the RE algorithm. It is easy to understand the shortcomings of the Ostrich algorithm in this application. Censoring reduces the magnitude of the more extreme movements in observed price and must dampen the fluctuations due to the periodicity present in the conditional variance of the equilibrium prices.

Table 9 : Treasury bill futures model estimates

$$y_t = \mu + \varepsilon_t$$
$$\varepsilon_t | I_{t-1} \sim N(0, \sigma_t^2)$$
$$\sigma_{t+1}^2 = \alpha_0 + \alpha_1 \varepsilon_t^2 + \alpha_2 \sigma_t^2 + \sum_{j=2,3,6} \gamma_j [D_{jt+1} - (\alpha_1 + \alpha_2) D_{jt}].$$

A. RE algorithm, limit moves treated as censored observations

μ	α_0	α_1	α_2	γ_2	γ_3	γ_6
0.0128	0.0199	0.042	0.927	-0.210	-0.276	0.449
(0.0335)	(0.0034)	(0.008)	(0.002)	(0.097)	(0.071)	(0.208)

B. Ostrich algorithm, limit move prices treated as equilibrium prices

μ	α_0	α_1	α_2	γ_2	γ_3	γ_6
0.0081	0.0117	0.040	0.939	-0.131	-0.229	0.284
(0.0312)	(0.0056)	(0.012)	(0.026)	(0.095)	(0.089)	(0.154)

Notes. Robust standard errors are in parentheses [Bollerslev and Wooldridge (1992)]. (γ_j, *j=2,3,6*) are the coefficients for indicator variables corresponding to Tuesday, Wednesday, and the first trading day following the H6 release of the Federal Reserve system, respectively. The dependent variable is scaled by multiplication by *1000*. On the days with limit moves of *0.60* in the IMM index value the median index value was approximately *87* and the corresponding rate of change of futures price was *2.07* times the unconditional standard deviation for the period.

Table 10 : Diagnostic tests of Treasury bill futures model

	RE	Ostrich
Residual analysis		
Z	0.47 (0.641)	0.47 (0.641)
Q(10)	4.73 (0.909)	4.94 (0.895)
$Q^2(10)$	9.31 (0.503)	10.28 (0.417)
Robust conditional moment tests for omitted variables: mean		
ε_{t-1}	0.69 (0.407)	1.17 (0.280)
D_1	2.74 (0.098)	2.55 (0.111)
D_2	0.15 (0.694)	0.22 (0.642)
D_3	0.57 (0.452)	0.50 (0.480)
D_4	0.75 (0.386)	0.70 (0.402)
D_5	0.77 (0.381)	0.71 (0.399)
Robust conditional moment tests for omitted variables: variance		
D_1	0.11 (0.739)	0.02 (0.898)
D_4	0.00 (0.962)	0.00 (0.984)
D_5	0.11 (0.745)	0.08 (0.780)

Notes: *p*-values are shown in parentheses. *Z* is a runs test statistic [Lehmann (1976, page 313)] with *p*-values from the unit normal distribution. *Q(10)* and $Q^2(10)$ are the Ljung Box (1978) portmanteau tests for autocorrelation in the standardized residuals and their squares; *p*-values are for the χ^2_{10} distribution. For the robust conditional moment tests the *p*-values are from the χ^2_{10} distribution. D_j, *j=1,...,5* are indicator variables for the day of the week.

5. Conclusion

Financial markets such as equity and futures markets often adopt rules for interrupting trading when sharp movements in price occur. Traders must then estimate the unobservable equilibrium price from the observed censored price. We have developed an algorithm for estimating GARCH models for markets subject to such price limits. The algorithm extends existing statistical methods for dealing with censoring to applications with heteroscedastic data. The result is a relatively simple procedure for accommodating these two important features of some futures and equity data. The RE algorithm replaces unobservable components of the likelihood function of the GARCH model by their rational expectations. It may be applied to a wide variety of censoring situations, whether the price limits are daily or intradaily, fixed or escalating.

We showed that our RE algorithm is an EM algorithm in situations where the imposition of price limits impacts the equilibrium price process such that its conditional variance is a function of observable data. In such cases it produces maximum likelihood estimates. In other cases it only approximates the maximum likelihood estimates, where the degree of approximation depends on its ability to track the true conditional variance process. The results of our Monte Carlo analysis suggest that in most situations the error of approximation is likely to be small--much smaller than that provided by the commonly used Ostrich algorithm which treats the censored data as equilibrium prices.

In an application to Treasury bill futures we found that, relative to our RE algorithm, the Ostrich algorithm tended to underestimate the periodicity in the conditional variance due to news such as the H6 announcements of the Federal Reserve System.

Promising areas for adoption of the RE algorithm include applications such as the pricing of options on futures, where accurate estimates of conditional variance are crucial.

References

Amin, K.I., and V.K. Ng, (1994), A comparison of predictable volatility models using option data, manuscript.

Baillie, R.T. and T. Bollerslev, (1989), The message in daily exchange rates: a conditional variance tale, *Journal of Business and Economic Statistics,* 7, 297-305.

Blight, B.J.N., (1970), Estimation from a censored sample for the exponential family, *Biometrika,* 57, 389-395.

Bollerslev, T., (1986), Generalized autoregressive conditional heteroscedasticity, *Journal of Econometrics,* 31, 307-327.

Bollerslev, T. and J.M. Wooldridge, (1992), Quasi-maximum likelihood estimation of dynamic models with time-varying covariances, *Econometric Reviews,* 11, 143-172.

Brennan, M.J., (1986), A theory of price limits in futures markets, *Journal of Financial Economics,* 16, 213-233.

Calzolari, G. and G. Fiorentini. (1997), A Tobit model with GARCH errors, *Econometric Reviews,* forthcoming.

Davidson, R., and J.G. MacKinnon, (1993), *Estimation and inference in econometrics,* [Oxford University Press, New York]

Dempster, A.P., N.M. Laird and D.R. Rubin, (1977), Maximum likelihood from incomplete data via the EM algorithm, *Journal of the Royal Statistical Society*, Series B, 39, 1-22.

Engle, R.F., (1982), Autoregressive conditional heteroscedasticity with estimates of the variance of United Kingdom inflation, *Econometrica,* 50, 987-1007.

Engle, R.F., and V.K. Ng, (1993), Measuring and testing the impact of news on volatility, *Journal of Finance,* 48, 1749-1801.

Hasselblad, V., A.G. Stead, and W. Galke, (1980), Analysis of coarsely grouped data from the lognormal distribution, *Journal of the American Statistical Association,* 75, 771-778.

Kodres, L.E., (1988), Tests of unbiasedness in foreign exchange futures markets: The effects of price limits, *Review of Futures Markets,* 7, 138-166.

Kodres, L.E., (1993), Tests of unbiasedness in the foreign exchange futures markets: An examination of price limits and conditional heteroscedasticity, *Journal of Business,* 66, 463-490.

Lehmann, E.L., (1975), *Nonparametrics: Statistical methods based on ranks.* (Holden-Day, San Francisco).

Ljung, G.M. and G.E.P. Box, (1978), On a measure of lack of fit in time series models, *Biometrika,* 65, 297-303.

Lumsdaine, R.L., (1995), Finite sample properties of the maximum likelihood estimator in GARCH(1,1) and IGARCH(1,1) models: A Monte Carlo investigation, *Journal of Business and Economics Statistics,* 13, 1-10.

McCurdy, T. and I.G. Morgan, (1987), Tests of the martingale hypothesis for foreign currency futures with time-varying volatility, *International Journal of Forecasting,* 3, 131-148.

McGill, J.I., (1995), Censored regression analysis of multiclass passenger demand data subject to joint capacity constraints, *Annals of Operations Research,* 60, 209-240.

McLeod, A.J. and W.K. Li, (1983), Diagnostic checking ARMA time series models using squared-residual autocorrelations, *Journal of Time Series Analysis,* 4, 269-273.

Maddala, G.S. (1983), *Limited dependent and qualitative variables in econometrics.* Economic Society Monographs, No. 3. (Cambridge University Press, New York).

Phillips, G., and R.Weiner, (1996), Does futures trading increase price volatility? Evidence from crude oil derivatives, manuscript.

Raiffa, H. and R. Schlaifer, (1961), *Applied statistical decision theory.* (Harvard University, Boston).

Trevor, R.G, (1994), GARCH with missing observations: A pseudo EM algorithm, manuscript.

Wooldridge, J.M., (1990), A unified approach to robust, regression-based specification tests, *Econometric Theory,* 6, 17-43.