Journal of Empirical Finance 16 (2009) 337–351

Contents lists available at ScienceDirect

Journal of Empirical Finance

journal homepage: www.elsevier.com/locate/jempfin

A censored stochastic volatility approach to the estimation of price limit moves

Ping-Hung Hsieh [1,2], J. Jimmy Yang [*,2]

Oregon State University, College of Business, 200 Bexell Hall, Corvallis, OR 97331-2603, USA

ARTICLE INFO

Article history:
Received 13 December 2006
Received in revised form 4 November 2008
Accepted 11 November 2008
Available online 21 November 2008

JEL classification:
C5
G1

Keywords:
Price limits
Stochastic volatility
Stock markets
Futures markets

ABSTRACT

A censored stochastic volatility model is developed to reconstruct a return series censored by price limits, one popular form of market stabilization mechanisms. When price limits are reached, the observed prices are truncated and the equilibrium prices are unobservable, which makes further financial analyses difficult. The model offers theoretically sound estimates of censored returns and is demonstrated via simulations to outperform existing approaches with respect to the estimates of model parameters, unconditional means, and standard deviations. The algorithm is applied to model stock and futures returns and results are consistent with the simulation outcomes.

1. Introduction

One of the major challenges faced by all empirical price limit researchers is the appropriate handling of price limit moves. In essence, none of the empirical studies that examine financial data subject to price limits are immune from this challenging task. When price limits are hit, the observed prices are truncated and the equilibrium prices are unobservable, which makes it difficult to calculate true returns and volatilities for further financial analyses.[3] Existing studies do not provide a generally accepted solution. We contribute to the literature by developing a censored stochastic volatility (CSV) model to capture important features of a return series censored by price limits.

Price limits are imposed in many financial markets to prevent excessive price fluctuation.[4] Although price limit rules vary from country to country and market to market, broadly speaking, they may be defined as boundaries set by market regulators to restrict daily security price changes. Because prices are constrained, well-established financial analytical methods and models, such as a mean-variance analysis or GARCH (generalized autoregressive conditional heteroscedasticity) modeling, may create biased results. In their comprehensive literature review, Kim and Yang (2004) conclude that it is necessary for future studies to develop models that might overcome potential biases in the statistical inference that are due to the constrained return-generating process.

* Corresponding author. Tel.: +1 541 737 6005; fax: +1 541 737 4890.
E-mail addresses: hsiehph@bus.oregonstate.edu (P.-H. Hsieh), jimmy.yang@bus.oregonstate.edu (J.J. Yang).
[1] Tel.: +1 541 737 4297; fax: +1 541 737 4890.
[2] Both authors thank participants at the 2004 Joint Statistical Meetings for their valuable comments on an earlier version of the paper. The helpful suggestions from two anonymous referees and the editor Franz Palm are greatly acknowledged. The usual disclaimer applies.
[3] The term “equilibrium price” refers to a price that is cleared by the market. When price limits are reached, the observed prices are not equilibrium prices because the market is not cleared due to likely order imbalance.
[4] For example, the U.S. futures markets and many stock exchanges around the world—including China, Japan, Korea, Malaysia, Taiwan, and Thailand in Asia and Austria, Belgium, France, Greece, Italy, the Netherlands, Spain, and Switzerland in Europe—currently adopt price limits.

0927-5398/$ – see front matter © 2008 Elsevier B.V. All rights reserved.
doi:10.1016/j.jempfin.2008.11.002

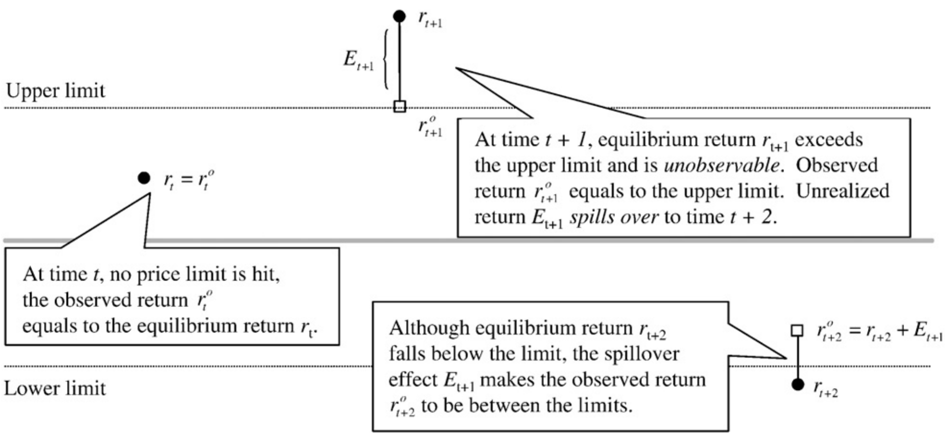

Fig. 1. Unobservability of price-limited data and the spillover effect.

Price-limited data share similar characteristics—such as heteroscedasticity and volatility clustering—with regular asset return data, but differ in two important ways. First, equilibrium prices are unobservable when price limits are hit, and observable prices are truncated at the limit. Second, when price limits are hit, the return beyond price limits must be reflected on the next trading day; thus, the unrealized return spills over to the following day, creating the spillover effect.[5] Fig. 1 provides a simple illustration of these two distinct features. At time t, the observed return r^o_t equals the equilibrium return r_t because it is within the price limit boundary.[6] At the next time period $t+1$, the equilibrium return r_{t+1} exceeds the upper price limit, but due to the price limit mechanism, we can only observe r^o_{t+1}. Because of the price limit constraint, the amount of the unrealized return E_{t+1} spills over to time $t+2$. Hence, though the equilibrium return r_{t+2} at time $t+2$ falls below the lower limit, the observed return r^o_{t+2} lies between the limits if the spillover portion E_{t+1} is included.

Our CSV model accounts for the unobservability and spillover effect associated with price limits.[7] Although several models have been developed to recognize the unobservable feature of price-limited data, to our knowledge, only Wei's (2002) censored-GARCH model has incorporated the spillover effect. However, as Wei (2002) acknowledges, the disadvantage of the censored-GARCH model is the long running time of its procedure, which makes it impractical to design a Monte Carlo simulation to verify its claimed performance. We are able to run simulations to show the superior performance of our CSV model because of the simple conditional likelihood function of its parameters and the resulting efficient algorithm for parameter estimation. We use a theoretically sound imputation method to generate fill-ins for censored observations and thereby create an imputed return series that is appropriate for subsequent financial analyses, such as option pricing or asset allocations. As suggested by Jacquier et al. (1994) and Andersen et al. (1999), we use the Markov chain Monte Carlo (MCMC) technique for efficient parameter estimation. Because our model addresses the spillover effect explicitly, a complete return series can be imputed even if price limits are reached on consecutive days, a situation that occurs in practice but cannot be handled appropriately by existing models.

The stochastic volatility (SV) models have recently gained popularity in modeling financial asset returns. Durham (2007) shows that SV models do a good job capturing the dynamic of volatility and the shape of the conditional distribution of financial asset returns. Nardari and Scruggs (2007) also utilize the SV process to capture the time-varying covariance matrix of returns in evaluating empirical implications of asset pricing theory. Our CSV model provides a realistic and flexible modeling of financial time series data in that it involves two noise processes and allows error terms to be correlated. Given the increasing adoption of SV models in empirical financial studies, our CSV model provides researchers an appropriate and efficient approach to handle price-limited data.

We compare our model with several existing approaches through a Monte Carlo simulation. The results suggest that our model outperforms other approaches with respect to the estimation of model parameters, the unconditional means, and the standard deviations. When the price limits are set at ±7%, our CSV model recovers censored returns and gives an estimate of standard deviation with less than 1% error. However, the standard deviation is underestimated by 5% when observed prices are used to calculate returns and by 14% when limited prices are deleted from the sample. The consequences of such underestimation can be substantial. First, an underestimation of volatility generally translates into the underpricing of an option, which can be costly to option traders. Second, the underestimation of risk affects asset allocation decisions by portfolio managers and capital budgeting decisions by corporate managers. In addition, results from empirical studies that use the underestimated standard deviation are likely to be biased.

We demonstrate the usefulness of our algorithm by modeling the returns of two actively traded stocks on the Taiwan Stock Exchange (TSE) and two U.S. futures contracts on the Chicago Board of Trade (CBOT) during volatile periods when price limit moves are more likely to occur. To examine the robustness of our results, we compare the performance of our CSV approach with that of

[5] Using an information-based argument, George and Hwang (1995) suggest that information is not fully reflected on the limit-hit day, and therefore, the remaining information must be reflected on the following day.

[6] The "equilibrium return" can be viewed as the true return when price limits were not present.

[7] Our model does not include the possible "magnet effect" proposed by Subrahmanyam (1994). Please see Berkman and Steenbeek (1998) and Cho et al. (2003) for models related to the magnet effect.

other existing models using the S&P 500 Index returns and the GARCH-generated returns. Overall, our result that the CSV approach outperforms other existing models is robust.

We organize the remainder of this article as follows: In the next section, we review existing methods of handling and modeling price-limited data in the literature. In Section 3, we describe our model and estimation approach. Then in Section 4, we present the simulation design and results, followed by some empirical results for stock and futures returns in Section 5. We perform robustness check in Section 6 and conclude in Section 7.

2. Handling and modeling price-limited data

Existing studies have applied both the ostrich (also known as the "naïve") and the discarding approaches to deal with price limit moves (e.g., Cho et al., 2003; Kim and Rhee, 1997; McCurdy and Morgan, 1987). The ostrich approach ignores price limits and treats the observed price as if it were actually the equilibrium price, while the discarding approach drops limited prices from the sample. Although the ostrich and discarding approaches are easy to implement, they undoubtedly generate misleading results, which in turn alter important financial factors and decisions. For example, discarding all price limit moves from the sample breaks down the dynamic structure of a price series; this approach becomes even more problematic when consecutive limit moves occur (Miller, 1989). Several studies demonstrated that results are biased when the ostrich or discarding approach is used to measure risk and returns under price limits (e.g., Chiang et al., 1990; Lee and Kim, 1997).

Several studies have attempted to provide some remedy for handling price-limited data. In his early study, Roll (1984) uses the closing price on the subsequent day without a limit move as the equilibrium price on the limit day. This approach breaks down the dynamic structure of a price series. Miller (1989) suggests sampling over a longer period to reduce the effects of limits; however, this method reduces the sample size substantially and thus decreases the statistical precision. Although Chou's (1997) Bayesian approach eliminates the need for large samples, the normality assumption of the model errors seems restrictive. Other methods that rely on the arbitrage theory (Kodres, 1993) and the price process itself (Sutrick, 1993) to handle price limit moves are numerically complex and constrained by theoretical assumptions.

In addition, several models have been developed to recognize the unobservable features of price-limited data, including censored regression (Kodres, 1988, 1993) and Tobit (Lee, 1999; Wei, 1999) models. Tobit-GARCH models (Calzolari and Fiorentini, 1998; Morgan and Trevor, 1999) further consider the heteroscedasticity feature of asset returns, but they do not address the spillover effect. Wei (2002) develops a censored-GARCH model to incorporate the spillover effect, but the complexity of the model demands a lengthy estimation process to obtain parameter estimates, especially when more parameters and/or price limit moves appear. Wei and Chiang (2004) derive a generalized method of moments (GMM) estimator of variances in markets with price limits to overcome the measurement problem of deleting prices on limit days. They find that the GMM is easier to implement, but it is difficult to extend its use to the case of conditional heteroscedasticity. Chung and Gan (2005) apply Wei and Chiang's (2004) method to examine whether price limits affect underlying equilibrium prices on limit days and show that the result is sensitive to the density of the return series they assume.

In sum, existing approaches and models are either unable to account for both the unobservability and the spillover effect associated with price limits or too computationally complicated to be feasible for simulation studies. This paper contributes to the literature by proposing the CSV model that overcomes problems faced by existing models.

3. The censored stochastic volatility model

In this section, we introduce the proposed model and describe a Bayesian approach to estimating the model parameters.

3.1. Stochastic volatility model with correlated errors

Recall from Fig. 1 that r_t^o and r_t are the continuously compounded observed and equilibrium log returns of the asset, respectively. Due to the price limit rules and the spillover effect, the relationship between r_t^o and r_t can be summarized as follows:

$$r_t^o = \begin{cases} u_t & \text{if } r_t + E_{t-1} \geq u_t \\ r_t + E_{t-1} & \text{if } d_t < r_t + E_{t-1} < u_t \, , \\ d_t & \text{if } r_t + E_{t-1} \leq d_t \end{cases} \tag{1}$$

where E_{t-1} is the unrealized return at the previous time period $t-1$ spilled over to time t, and u_t and d_t are the upper and lower return limits derived from the price limit rules. E_{t-1} directly measures the spillover effect. If price limits are not reached at time $t-1$, there is no spillover effect and E_{t-1} is equal to zero as there is no unrealized return. Eq. (1) recaps our discussion of Fig. 1, according to which if the added return (i.e., the sum of the equilibrium return r_t and the spillover effect E_{t-1}) hits the limit u_t (or d_t), the observed return r_t^o equals the limit u_t (or d_t); otherwise, r_t^o equals the sum of r_t and E_{t-1}. Note that the return limits u_t and d_t may vary or be constant depending on the price limit rules.[8] The relationship between the equilibrium and observed returns and prices can be derived through simple algebra (e.g., Eqs. (1) and (2) in Wei, 2002).

[8] For example, the price limits on the U.S. futures markets are contract specific and based on fixed dollar amounts, whereas most stock markets impose percentage-based price limit rules. That is, u_t and d_t may vary from day to day in the U.S. futures markets, but they are constant for percentage-based limits. Section 5, application to stock and futures returns, provides more details.

Following the argument in the literature that price limits only affect the observability of the equilibrium prices, not their underlying processes (see, e.g., Wei, 2002; Wei and Chiang, 2004), we propose the SV model (Tauchen and Pitts, 1983; Taylor, 1982) to model the process of the equilibrium return r_t. Specifically, the model in its canonical form is

$$r_t = e^{h_t/2}\varepsilon_t, \text{ and} \tag{2}$$

$$h_t = \mu + \phi(h_{t-1} - \mu) + \sigma_v v_t, \tag{3}$$

where $-1<\phi<1$, and, for a basic stochastic volatility model, ε_t and v_t for $t=1,\ldots,n$ are independent random variables from the standard normal distribution. As both the asset returns and their conditional volatility are derived from asset prices, the shocks to asset returns and shocks to asset volatility could be correlated. We extend the basic model to allow the correlation between ε_t and v_t with a correlation parameter ρ. That is, ε_t and $\sigma_v v_t$ are jointly normal with a mean vector $[0\ 0]'$ and the covariance matrix $\Sigma = \begin{pmatrix} 1 & \rho\sigma_v \\ \rho\sigma_v & \sigma_v^2 \end{pmatrix} = \begin{pmatrix} 1 & \psi \\ \psi & \Omega + \psi^2 \end{pmatrix}$, where the parameters ψ and Ω are introduced to increase the efficiency in parameter estimation, see Appendix A and Jacquier et al. (2004). Note that h_t is referred to as the latent log volatility and measures the amount of volatility on trading day t; ϕ, measuring the autocorrelation present in the logged squared data, can be interpreted as the persistence in volatility. Finally, σ_v is considered the conditional volatility of log volatilities (Kim et al., 1998; Meyer and Yu, 2000).

3.2. Chained data augmentation

In this section, we discuss a simulation-based parameter estimation algorithm. We denote the set of model parameters and log volatilities as $\theta = \{\mu, \phi, \Omega, \psi, h_0\} \cup \{h_t, t = 1, \ldots, n\}$, the set of observed returns as $R^o = \{r_1^o, r_2^o, \ldots, r_n^o\}$, and a set of indicators as $\Delta = \{\delta_1, \delta_2, \ldots, \delta_n\}$, where $\delta_t = 1$ (or -1) if the upper (lower) limit is hit at time t and 0 otherwise. With the set Δ, we define $T_u = \{t: \delta_t = 1, 1 \le t \le n\}$ and $T_d = \{t: \delta_t = -1, 1 \le t \le n\}$ as the sets of time indices during which the upper and lower limits are hit, respectively. Similarly, $T_o = \{t: \delta_t = 0, 1 \le t \le n\}$ contains a set of time indices when neither limit is hit. The set of unobserved equilibrium returns is defined as R^*.

Consider the following two equations:

$$p(\theta|R^o) = \int_{R^*} p\left(\theta|R^o, R^*\right) \cdot p\left(R^*|R^o\right) dR^*, \text{ and} \tag{4}$$

$$p\left(R^*|R^o\right) = \int_{\theta} p\left(R^*|R^o, \theta\right) \cdot p(\theta|R^o) d\theta, \tag{5}$$

where $p(\theta|R^o)$ is the posterior density of the model parameter θ given the observed return R^o, $p(R^*|R^o)$ is the predictive density of the latent returns R^* given the observed returns R^o, $p(\theta|R^o,R^*)$ is the conditional density of θ given the augmented returns R^o and R^* (i.e., the augmented posterior), and $p(R^*|R^o,\theta)$ is the conditional predictive density. By applying this method of composition to Eqs. (4) and (5), we obtain the following iterative algorithm[9]: At the ith iteration,

- Step A. Generate $\theta_{(i)}$ from $p(\theta|R^o, R^*_{(i-1)})$.
- Step B. Generate $R^*_{(i)}$ from $p(R^*|R^o, \theta_{(i)})$.

The algorithm iterates and generates both $\theta_{(i)}$ and $R^*_{(i)}$ for $i=1,2,\ldots$. Following the Markovian properties (Tanner and Wong, 1987), we have

$$\theta_{(i)} \xrightarrow{d} \theta \sim p(\theta|R^o), \text{ and} \tag{6}$$

$$\frac{1}{m'} \sum_{i=m+1}^{m+m'} f\left(\theta_{(i)}|R^o\right) \xrightarrow{\text{a.s.}} \mathrm{E}[f(\theta)|R^o], \tag{7}$$

for any integrable function f.

The implementation of the iterated algorithm (Steps A and B) relies on drawing samples from the conditional distributions of $p(\theta|R^o,R^*_{(i-1)})$ and $p(R^*|R^o,\theta_{(i)})$. To generate random variates from non-standard distributions, we employ numerical techniques described in Kim et al. (1998) and include the discussion of Steps A and B in Appendix A.[10]

4. Performance assessment

In this section, we explore the performance of the CSV by comparing its estimation errors to those from the existing models with respect to estimates of model parameters, unconditional means, and standard deviations.

[9] For a detailed derivation, see Tanner (1993).

[10] Other numerical methods that deal with non-standard distributions can be found in Ritter and Tanner (1992), Gilks and Wild (1992), Neal (1997), Gilks et al. (1998). Detailed descriptions of the Steps A and B are available on request.

Table 1
Simulation results

Panel A

N	Bounds	# of limit hits	# of upper limit hits	# of lower limit hits
100	±7%	5.55	2.63	2.92

Panel B

Variable	Ostrich		Discarding		CSV	
	RMSE ($\times 10^{-4}$)	MPE (%)	RMSE ($\times 10^{-4}$)	MPE (%)	RMSE ($\times 10^{-4}$)	MPE (%)
μ	1667	1.5	3196	4.3	1527	0.1
ϕ	1207	−12.7	1324	−14.2	1069	−9.8
σ_v	134	−1.0	158	−0.4	142	1.9
ρ	979	17.0	1099	17.4	945	8.2

Panel C

Method	Unconditional mean		Unconditional standard deviation	
	RMSE ($\times 10^{-4}$)	MPE (%)	RMSE ($\times 10^{-4}$)	MPE (%)
Ostrich	4.68	−8.6	21.93	−4.8
Discarding	20.33	−13.9	56.90	−13.6
CSV	4.98	−4.1	15.01	0.8
GMM	4.68	−8.6	29.21	−5.4

This table reports the simulation results from the following stochastic volatility model:

$$r_t = e^{h_t/2}\varepsilon_t, \text{ and}$$

$$h_t = \mu + \phi(h_{t-1} - \mu) + \sigma_v v_t$$

where ε_t and $\sigma_v v_t$ are jointly normal with a mean vector $[0\ 0]'$ and the covariance matrix $\Sigma = \begin{pmatrix} 1 & \rho\sigma_v \\ \rho\sigma_v & \sigma_v^2 \end{pmatrix}$ for $t=1,\ldots,n$, and $-1<\phi<1$. ρ is the correlation between ε_t and v_t. h_t is referred to as the latent log volatility that measures the amount of volatility on trading day t, and ϕ, which measures the autocorrelation present in the logged squared data, can be interpreted as the persistence in the volatility. In addition, σ_v is considered the volatility of log-volatilities. The root mean squared error $(\text{RMSE}) = \sqrt{\frac{1}{m}\sum_{i=1}^{m}\left(\hat{\theta}_i - \theta\right)^2}$, and the mean percentage error $(\text{MPE}) = \frac{1}{m}\sum_{i=1}^{m}\left(\frac{\hat{\theta}_i - \theta}{\theta}\right)$, where $\hat{\theta}_i$ is either a parameter estimate (Panel B) or sample mean/standard deviation estimates (Panel C) of the ith sample and $m=1000$ in the MCMC simulation, θ is the corresponding true value. The CSV is our proposed censored stochastic volatility model, the Ostrich approach ignores price limits, and the discarding approach deletes limit-hit observations. Panel A reports the average number of limit hits from 1000 simulated return series. Bounds are the price limit rates. Upper limit hits occur when the price hits the upper bound, and lower limit hits occur when the price hits the lower bond. Limit hits include both the upper and lower limit hits. n is the sample size. Panel B reports the estimations of model parameters and Panel C reports the unconditional mean and standard deviation from our MCMC simulation. Posterior standard deviations are in parentheses. The generalized method of moments (GMM) estimator is that proposed by Wei and Chiang (2004).

4.1. Simulation design

Starting with an unlimited return (i.e., $r_1 = r_1^o$ and thus $E_1 = 0$), we calculate the "observed" returns r_t^o, $t \geq 2$ sequentially using the definition in Eq. (1). That is, r_t^o, $t \geq 2$ is set to the upper (lower) limit if the sum of the simulated equilibrium return r_t and the spillover effect from the previous time period E_{t-1} exceeds the upper (lower) limit. Otherwise, $r_t^o = r_t + E_{t-1}$. In our simulation, we set the model parameters as $\phi=0.8$, $\mu=-6.5$, and $\sigma_v=0.15$ on the basis of a real asset return series we analyze and experiment with various values of ρ. Results of $\rho=0.3$ are reported.

To increase simulation efficiency, we generate 100 samples of a size of 100 using the parameters defined above to select appropriate hyperparameter values. The parameter estimates from these "training" samples lead to the following prior distribution: $\mu \sim N(-6.5, 1)$, $\phi \sim Beta(10, 1.5)$, $\Omega^{-1} \sim Gamma(0.11, 7.15)$ and $\psi | \Omega \sim N(0.0, \Omega/2)$. The MCMC sampler is initialized by setting μ and h_t, $t=0,2,\ldots,n$ to the logarithm of variance of the sample, $\phi=\rho=0.5$, and $\sigma_v^2=0.35$. All upper (lower) limit hits are set to the upper (lower) bound to initialize the process.

We generate 1000 samples of size 100. For each sample, we run 251,000 iterations and retain the results from the last 1000 iterations for analysis. The long chain required for convergence is understandable given the number of parameters that needed to be estimated. We evaluate the performance of our CSV against the ostrich, discarding, and GMM (Wei and Chiang, 2004) methods. For the ostrich approach, we treat each sample as a realization from a basic SV model with correlated errors and apply the MCMC algorithm to estimate the parameters, using the same priors and initial values. We repeat this process after removing the limited returns for the discarding approach.

4.2. Simulation results

We report the simulation results in Table 1.[11] When price limits are set at ±7%, on average 5.55 (2.63 upper and 2.92 lower) limit hits occur in a sample of 100 as reported in Panel A. Following the Markov chain property (Eq. (7)), we average the last 1000 draws

[11] We only report the results for the sample size of 100 and price limits of ±7% due to space limitation. Results for the sample size of 250 and price limits of ±9% and ±5% are available on request. In general, the larger the sample size and the narrower the price limits, the better the CSV performs relative to other methods as more limit hits are observed.

of $\theta \sim p(\theta|R^o,R^*)$ in the Markov chain to obtain the posterior mean as an estimate of θ. For all three models, we evaluate the estimation results with the root mean squared error (RMSE),

$$\sqrt{\frac{1}{m}\sum_{i=1}^{m}\left(\hat{\theta}_i-\theta\right)^2}, \tag{8}$$

and mean percentage error (MPE),

$$\frac{1}{m}\sum_{i=1}^{m}\left(\frac{\hat{\theta}_i-\theta}{\theta}\right), \tag{9}$$

where $\hat{\theta}_i$ is a parameter estimate of the ith sample and $m=1000$.

In Panel B of Table 1, we summarize the performance of various methods. It is not surprising that the discarding approach has the worst performance in estimating model parameters. By removing all limited observations, the underlying dynamics have been altered; as a result, the model parameters cannot be reasonably estimated. Although the ostrich approach performs better than the discarding approach, our CSV appears to be the best in terms of the estimation error measured by RMSE and MPE.

A better estimate of the unconditional standard deviation from a limited return series is of great importance for various financial analyses; for example, it is one of the key ingredients in option pricing and mean variance analysis. We summarized the unconditional mean and standard deviation estimates in Panel C of Table 1. For the CSV, we imputed the limited returns $\{r_t^o : t \in T_u \cup T_d\}$ with the mean of the predictive distribution of $p(R^*|R^o,\theta)$ to create a fabricated sample. We then applied the conventional standard deviation formula to the fabricated sample to obtain an estimate of the unconditional standard deviation. For the ostrich and discarding approaches, unobserved returns are either substituted with the limits or discarded before the conventional standard deviation formula was applied to obtain an estimate. Finally, we included the GMM estimates, as proposed by Wei and Chiang (2004).

Both the ostrich and discarding approaches largely underestimate the standard deviation due to the substitution or removal of the unobserved returns. The problem gets worse when the number of limit hits increases. When the limits are set at ±7%, the ostrich and discarding approaches underestimate the standard deviation by 4.8% and 13.6%, respectively. The GMM does not seem to perform better than the other approaches as it underestimates the standard deviation by 5.4%.[12] As mentioned earlier, such underestimation can have significant impact on option pricing and financial decisions. The CSV provides substantial improvements in the estimation of standard deviation with only 0.8% error. The additional information provided by the CSV, including the persistence in volatility and theoretically sound imputed return values, allows insight into the price limit mechanism and the effects of the price limits.

5. Application to stock and futures returns

We demonstrate the implementation and performance of the CSV model by examining the trading data for two active stocks on the TSE and two futures contracts on the CBOT. Although price limits are present in both stock and futures markets, their impact may differ because of the different characteristics and regulation in these two types of markets. First, the price limit system in the TSE is based on a fixed percentage, whereas the CBOT sets a fixed dollar amount with possible expansion. Second, the difference in the degree of margin requirement, the mark-to-market feature in futures markets, and the market makers' obligation to provide liquidity in some stock markets may affect investors' trading behavior. Our empirical results show that the impact of price limits does differ in these two markets, but our model provides theoretically sound estimates of the true returns for both by recovering censored returns.

5.1. Stock market returns

The TSE is one of the most liquid and volatile markets in the world and has imposed daily price limits since its inception in 1962. The stated purpose of price limits is to avoid excessive volatility and protect investors by limiting their potential daily losses. The TSE sets both its upward and downward daily price limits at a predetermined rate on the basis of the previous day's closing price; this rate has been adjusted up or down several times in the past, as we report in Appendix B. Since 1989, the price limits have been set at 7%, with the downward limit being adjusted to 3.5% in accordance with the market conditions. Stocks that hit their price limits can be traded as long as the transaction prices are within the limits. Thus, the TSE price limits are simply boundaries, not triggers for trading halts.

To illustrate our estimation empirically, we use the stock returns of the Taiwan Semiconductor Manufacturing Company (TSMC) and Advanced Semiconductor Engineering (ASE). The TSMC is the world's largest dedicated semiconductor foundry, and ASE is the world's largest provider of independent semiconductor manufacturing services in assembly and testing. Both stocks are among the most actively traded in Taiwan. Furthermore, both companies have their American Depositary Receipts (ADRs) traded on the New York Stock Exchange; TSMC's ADR is one of the most active ADRs. Given the high trading activity and visibility of both stocks, we believe our results pertaining to them will be representative.

[12] It should be noted that the GMM estimates in our simulation are based on formulas that assume iid as reported in Wei and Chiang (2004).

Table 2
Empirical results for stock returns

Panel A

Stock	Number of trading days	Single-day limit hits		Consecutive limit hits		
		Upper	Lower	2 days	3 days	4 days
TSMC	841	26	9	3 (up, up) (down, down) (up, up)	1 (down, down, up)	1 (up, up, up, up)
ASE	841	27	15	3 (down, up) (down, down) (up, up)	1 (up, up, up)	2 (down, down, up, down) (up, up, up, up)

Panel B

Variable	TSMC			ASE		
	Ostrich	Discarding	CSV	Ostrich	Discarding	CSV
μ	−7.507 (0.325)	−7.730 (0.380)	−7.457 (0.270)	−7.246 (0.285)	−7.582 (0.263)	−7.182 (0.209)
ϕ	0.970 (0.012)	0.975 (0.012)	0.955 (0.016)	0.971 (0.013)	0.969 (0.014)	0.939 (0.023)
σ_v	0.209 (0.039)	0.164 (0.024)	0.304 (0.049)	0.165 (0.032)	0.140 (0.023)	0.314 (0.058)
ρ	0.091 (0.125)	0.001 (0.138)	*0.137 (0.098)	−0.067 (0.110)	−0.045 (0.148)	0.014 (0.097)

Panel C

Statistic	TSMC				ASE			
	Ostrich	Discarding	CSV	GMM	Ostrich	Discarding	CSV	GMM
Min	−7.203	−6.625	−10.916	N/A	−7.257	−6.811	−18.164	N/A
Mean	0.141	−0.014	0.197	0.141	0.148	0.032	0.180	0.148
Max	6.759	6.178	11.930	N/A	6.766	6.334	13.881	N/A
S.D.	2.769	2.349	3.072	3.042	2.944	2.450	3.359	3.112
Exc. Kurtosis	0.346	0.079	1.705	N/A	0.136	−0.075	2.315	N/A

*Significant at the 10% level.
Panel A reports the number of limit hits for TSMC and ASE. Panel B reports the estimation of parameters from the following stochastic volatility model using posterior means:

$$r_t = e^{h_t/2}\varepsilon_t, \text{ and}$$

$$h_t = \mu + \phi(h_{t-1} - \mu) + \sigma_v v_t$$

where ε_t and $\sigma_v v_t$ are jointly normal with a mean vector [0 0]′ and the covariance matrix $\Sigma = \begin{pmatrix} 1 & \rho\sigma_v \\ \rho\sigma_v & \sigma_v^2 \end{pmatrix}$ for $t = 1,\ldots,n$, and $-1 < \phi < 1$. ρ is the correlation between ε_t and v_t. h_t is referred to as the latent log volatility that measures the amount of volatility on trading day t, and ϕ measures the autocorrelation present in the logged squared data. σ_v is considered the volatility of log-volatilities. Posterior standard deviations are in parentheses. Panel C reports estimations of unconditional means and standard deviations (S.D.), and the minimum (Min) and maximum (Max) returns in percentage. Exc. Kurtosis is the excess kurtosis. The CSV is our proposed censored stochastic volatility model, the Ostrich approach ignores price limits, and the discarding approach deletes limit-hit observations. The generalized method of moments (GMM) estimator is that proposed by Wei and Chiang (2004).

We collected the daily closing prices, cash dividends, and stock dividends of both stocks from the TSE during January 4, 1996 to December 31, 1998 for a total of 841 trading days. Stock prices are adjusted for cash and stock dividends. It should be noted that few stocks hit the exact 7% price limit because of the tick sizes set by the TSE.[13] For a stock with a closing price of NT$80, the price on the next trading day may go up to NT$85.6 or down to NT$74.4. However, the tick size for the price range NT$50–150 is NT$0.5. Therefore, a limit hit occurs when the price rises to NT$85.5 or declines to NT$74.5, even though the return is only 6.875% or −6.875%, respectively.

Similar to our approach in Section 4, we generate 100 "training" samples in order to select appropriate hyperparameter values in the prior distributions. The parameter estimates from these "training" samples lead to the following prior distribution: $\mu \sim N(-9.0, 0.4)$, $\phi \sim Beta(10, 1.5)$, $\Omega^{-1} \sim Gamma(0.11, 7)$ and $\psi|\Omega \sim N(0.0, \Omega/2)$. The MCMC sampler was initialized by setting μ and h_t, $t = 0,2,\ldots,n$ to the logarithm of variance of log returns r_t^o, $\phi = \rho = 0.5$, and $\sigma_v^2 = 0.35$. The initial values of limited return r_t, $t \in R$ were set to the observed returns r_t^o and the corresponding E_t were set to 0. We run 251,000 iterations to minimize the impact of initial values. We summarize the results of the last 1000 iterations in Table 2.

Panel A of Table 2 reports the number of limit hits—when the closing price hits the limit—for both TSMC and ASE. During our sample period, TSMC experienced 35 (26 upper and 9 lower) single-day limit hits and 5 consecutive limit hits that ranged from two to a maximum of four days, whereas ASE had 42 (27 upper and 15 lower) single-day limit hits and 6 consecutive limit hits. Apparently, ASE hits the limits more often than TSMC, especially lower limits. The ratio of price limit days to the total number of observations is 5.7% for TSMC and 7.0% for ASE, consistent with our simulation results in Table 1. However, the empirical results demonstrate more upper than lower limit hits, an asymmetric property by default not observed in our simulation.

As we report in Table 2, more than 25% of the limit hits are associated with consecutive limit hits. In addition, these consecutive limit hits do not always move in the same direction. In our sample, three of the eleven consecutive limit hits are mixed upper and

[13] The tick sizes on the TSE are: NT$0.01 for stocks under NT$5, NT$0.05 for stocks between NT$5 and NT$15, NT$0.1 for stocks between NT$15 and NT$50, NT $0.5 for stocks between NT$50 and NT$150, NT$1 for stocks between NT$150 and NT$1000, and NT$5 for stocks NT$1000 and above, where NT$ is the Taiwanese currency.

lower limits. The empirical evidence reveals that contribution from any model that is unable to handle consecutive limit hits will be minimal. Our CSV model is capable of not only handling consecutive limit hits but also dealing with mixed consecutive limit hits effectively.

We report the parameter estimations using the posterior mean from the ostrich, discarding, and our CSV approaches for both TSMC and ASE in Panel B of Table 2. The ϕ parameter, which measures the autocorrelation in the logged squared data and can be interpreted as the persistence in the volatility, is the lowest for CSV. In contrast, σ_v, which is considered the volatility of log-volatilities, is the highest for our CSV. The correlation parameter ρ is positive and significant at the 10% level for the TSMC, but it is not significant for ASE. Our CSV is flexible in that it allows error terms ε_t and v_t to be correlated, so it works well for both the TSMC and ASE. We plot the estimated return series from the CSV approach in Fig. 2, in which the horizontal axis represents the trading day, and the vertical axis represents the daily return. With the 7% price limits, all daily returns are constrained, and thus, the observed returns are not equilibrium or true returns. As we show in Fig. 2, our CSV model is able to recover the censored returns and provide theoretically sound estimates of the true returns.

Panel C of Table 2 reports the estimation of the unconditional means, minimums, maximums, standard deviations, and excess kurtosis of log returns for the ostrich, discarding, CSV, and GMM approaches. Because all limit hits are discarded, the discarding approach has the highest minimum and the lowest maximum returns among all approaches. Consistent with our simulation results, our empirical evidence shows that the discarding approach has the lowest standard deviation because the most volatile

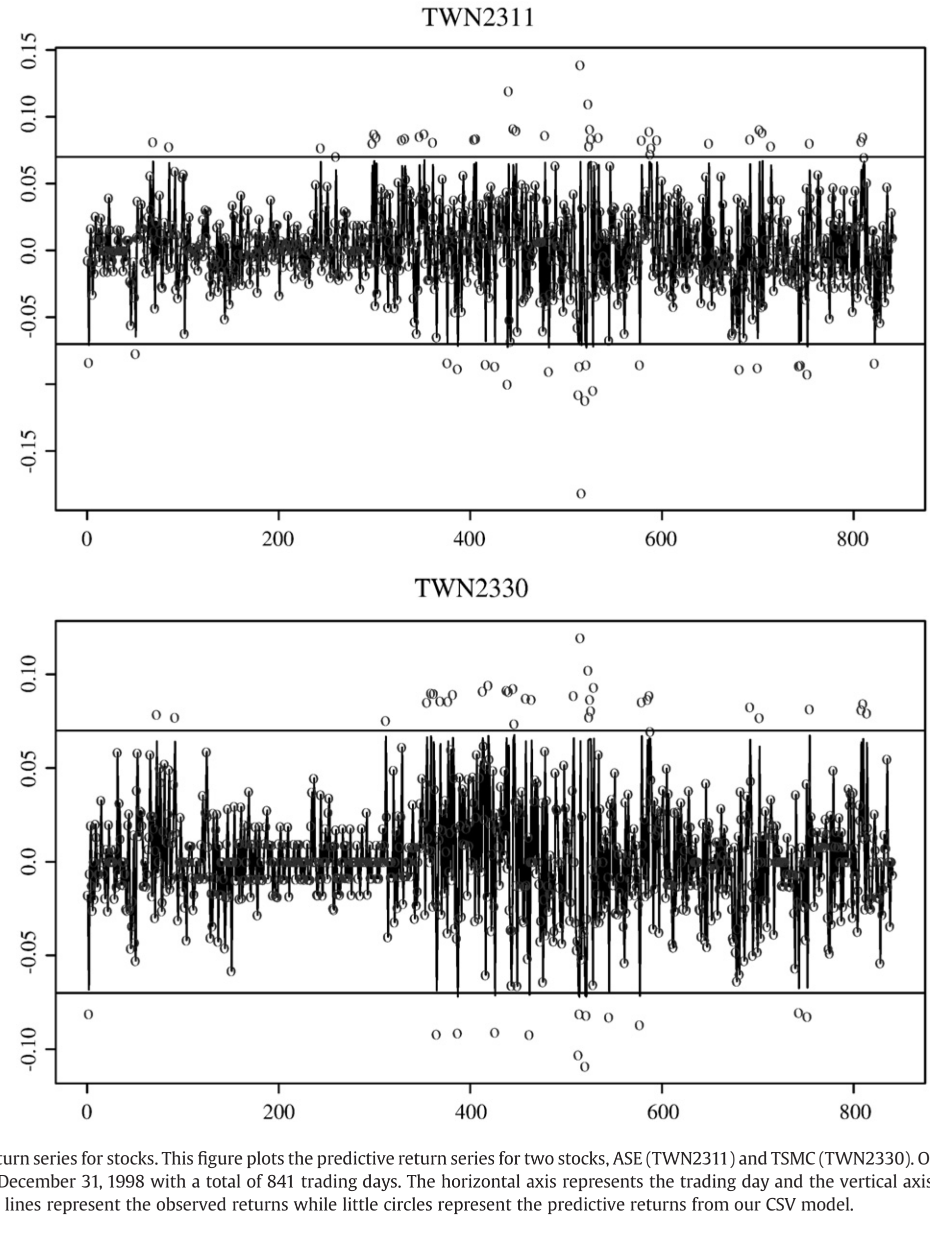

Fig. 2. Predictive return series for stocks. This figure plots the predictive return series for two stocks, ASE (TWN2311) and TSMC (TWN2330). Our sample spans from January 4, 1996 to December 31, 1998 with a total of 841 trading days. The horizontal axis represents the trading day and the vertical axis represents the daily return. The vertical lines represent the observed returns while little circles represent the predictive returns from our CSV model.

Table 3
Empirical results for futures returns

Panel A

Futures	Number of trading days	Single-day limit hits		Consecutive limit hits	
		Upper	Lower	2 days	3 days
Corn	649	3	4	0	0
Wheat	399	3	2	0	1 (down, down, up)

Panel B

Variable	Corn futures			Wheat futures		
	Ostrich	Discarding	CSV	Ostrich	Discarding	CSV
μ	−9.957 (0.965)	−9.961 (0.874)	−9.935 (1.003)	−8.756 (0.576)	−8.790 (0.470)	−8.787 (0.615)
ϕ	0.989 (0.007)	0.988 (0.007)	0.990 (0.006)	0.975 (0.015)	0.969 (0.019)	0.973 (0.016)
σ_v	0.197 (0.036)	0.195 (0.030)	0.193 (0.030)	0.164 (0.031)	0.159 (0.029)	0.185 (0.038)
ρ	0.204 (0.148)	0.132 (0.139)	0.201 (0.164)	−0.029 (0.193)	−0.064 (0.187)	−0.053 (0.156)

Panel C

Statistic	Corn futures				Wheat futures			
	Ostrich	Discarding	CSV	GMM	Ostrich	Discarding	CSV	GMM
Min	−3.451	−3.393	−4.450	N/A	−5.191	−4.837	−5.247	N/A
Mean	0.012	0.017	0.010	0.012	0.027	0.027	0.032	0.027
Max	3.577	3.295	4.552	N/A	5.191	4.831	7.378	N/A
S.D.	0.987	0.927	1.030	1.035	1.453	1.350	1.534	1.470
Exc. Kurtosis	2.265	1.899	3.816	N/A	1.178	0.849	2.529	N/A

Panel A reports the number of limit hits for the corn and wheat futures. Panel B reports the estimation of parameters from the following stochastic volatility model using posterior means:

$$r_t = e^{h_t/2}\varepsilon_t, \text{ and}$$

$$h_t = \mu + \phi(h_{t-1} - \mu) + \sigma_v v_t$$

where ε_t and $\sigma_v v_t$ are jointly normal with a mean vector $[0\ 0]'$ and the covariance matrix $\Sigma = \begin{pmatrix} 1 & \rho\sigma_v \\ \rho\sigma_v & \sigma_v^2 \end{pmatrix}$ for $t = 1,\ldots,n$, and $-1 < \phi < 1$. ρ is the correlation between ε_t and v_t. h_t is referred to as the latent log volatility that measures the amount of volatility on trading day t, and ϕ measures the autocorrelation present in the logged squared data. σ_v is considered the volatility of log-volatilities. Posterior standard deviations are in parentheses. Panel C reports estimations of unconditional means and standard deviations (S.D.), and the minimum (Min) and maximum (Max) returns in percentage. Exc. Kurtosis is the excess kurtosis. The CSV is our proposed censored stochastic volatility model, the Ostrich approach ignores price limits, and the discarding approach deletes limit-hit observations. The generalized method of moments (GMM) estimator is that proposed by Wei and Chiang (2004).

movements are deleted. The standard deviation of our CSV approach for TSMC (ASE) is 30.8% (37.1%) higher than that from the discarding, 10.9% (14.1%) higher than that from the ostrich, and 1.0% (7.9%) higher than that from the GMM. Overall, our empirical results provide support for our simulation analyses.

Price limits could affect the tail behavior of the return distribution, so it is possible that the excess kurtosis of the censored data becomes negative due to price limits. In other words, price limits may generate thin tails rather than fat tails of return distribution documented in the literature. Our result shows that the discarding approach leads to a negative excess kurtosis for ASE, consistent with our conjecture. Although the excess kurtosis for the ostrich approach is positive, the figure is relatively small. The CSV provides the highest excess kurtosis that is comparable to those reported in the literature (see e.g., Kon, 1984; Andersen et al., 2001). The ostrich and discarding approaches limit the occurrence of extreme movements and thus alter the tail behavior of the return distribution. Our CSV recovers the censored returns so the true fat-tailed distribution can be revealed.

Note that the minimum return for the ostrich approach is below the −7% lower limit because we use the log return but the TSE does not. In the previous example, a price drop from NT$80 to NT$74.5 would be considered −6.875% at the TSE, but the log return is actually −7.12%. In addition, the minimum and maximum for the GMM are not applicable because it groups returns on the basis of the different number of days, and thus, the minimum and maximum are not representative. Our CSV approach estimates the maximum daily return for TSMC (ASE) to be 11.93% (13.88%) and the minimum to be −10.92% (−18.16%).

5.2. Futures market returns

On October 15, 1935, the CBOT passed a resolution to impose price limits on futures contracts. The CBOT states that “price limits were established because of a perceived need for a *cooling off period* for futures when a major news release or event causes prices to move substantially.”[14] Unlike the TSE, the CBOT adopts a system that sets price limits symmetrically at a fixed dollar amount above and below the previous day's settlement price and in which no trading is allowed beyond the preset price range. Prior to

[14] Several studies have tried to test the effectiveness of price limits in futures markets. Please see Ma et al. (1989) and Chen (1998) for example or Kim and Yang (2004) for a more comprehensive survey.

August 27, 2000, price limits could be expanded after a contract closed at the limit price for consecutive trading days and then reverted back to the original level when price limits were not reached. However, since August 27, 2000, the CBOT has set wider price limits and removed the rule that expands price limits when limits are reached.

We use the trading data of corn and wheat futures contracts to illustrate the estimation of our model. Of all the contracts from 1960 to 2005, the 1996 contracts seem to be the most volatile and limit hits are more likely to occur. The trading of the December 1996 corn futures contracts began on May 31, 1994 and ended on December 19, 1996 for a total of 649 trading days; trading of the December 1996 wheat futures contracts began on May 25, 1995 and ended on December 19, 1996 for a total of 399 trading days. Panel A of Appendix C provides the price limits ($0.12 per bushel) for corn futures, and Panel B exhibits those ($0.20 per bushel) for wheat futures during the trading period. The scale of the price limit expansion is set at 50%, so the expanded price limit is $0.18 per bushel for corn futures and $0.30 per bushel for wheat futures. As the appendix shows, the expanded price limits lasted for two trading days for corn futures and three trading days for wheat futures before they reverted back to their normal levels.

We report the empirical results for corn and wheat futures in Table 3. As Panel A shows, the corn futures contract has seven single-day limit hits (three upper and four lower) and no consecutive limit hits, whereas the wheat futures contract has five single-day limit hits (three upper and two lower) and one consecutive limit hit that lasted for three days. The ratio of price limit days to the total number of observations is 1.1% for corn futures and 2.0% for wheat futures, significantly lower than the corresponding

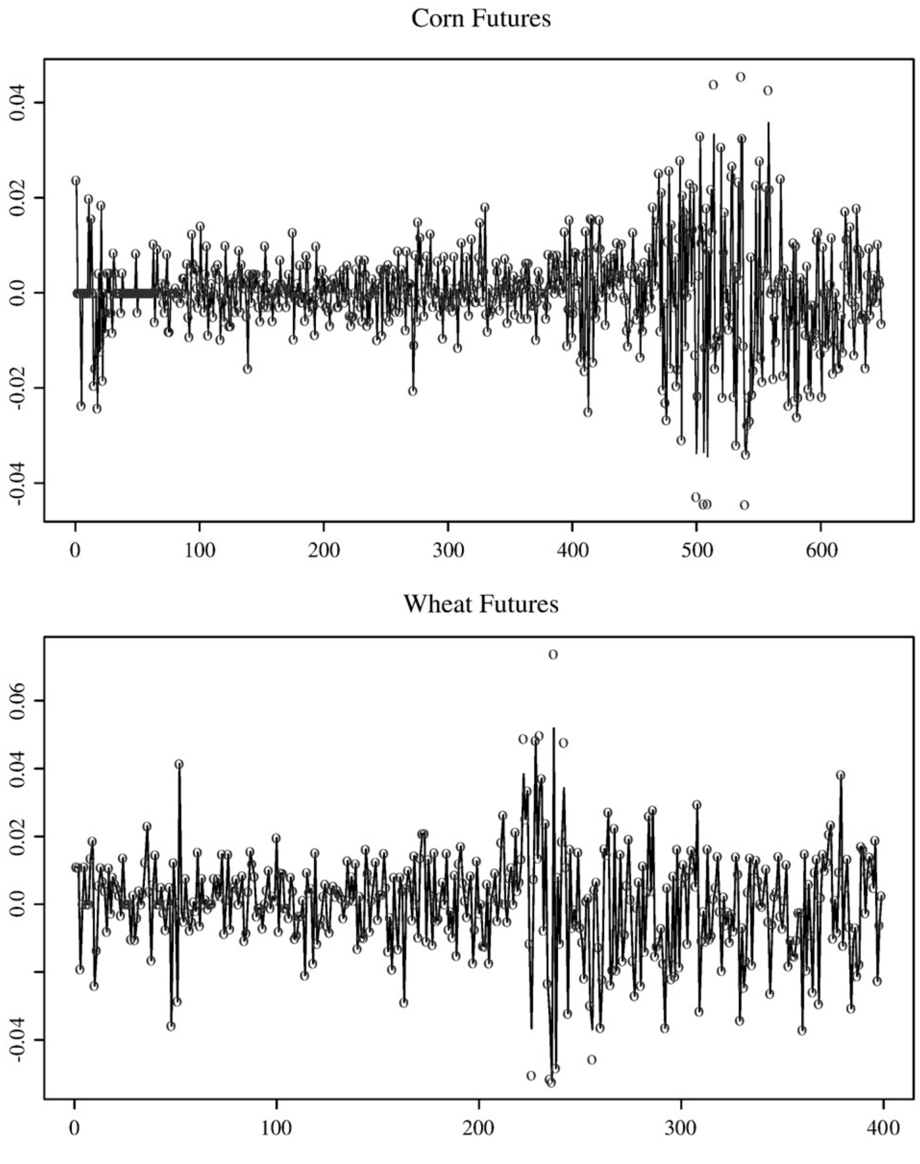

Fig. 3. Predictive return series for futures. This figure plots the predictive return series for corn and wheat futures. The trading of the corn futures contracts begins on May 31, 1994 and ends on December 19, 1996 for a total of 649 trading days; the trading of wheat futures contracts begins on May 25, 1995 and ends on December 19, 1996 for a total of 399 trading days. The horizontal axis represents the trading day, and the vertical axis represents the daily return. The vertical lines represent the observed returns, and the little circles indicate the predictive returns from our CSV model.

ratios for TSMC and ASE. However, upper and lower limit hits are more symmetrically distributed for futures than for stocks. These empirical results support our prior argument that the impact of price limits may differ because of the different characteristics and regulation in stock and futures markets.

Our CSV approach works on futures returns as well as on the stock returns. We report the parameter estimations using posterior mean from the ostrich, discarding, and our CSV approaches for both corn and wheat futures in Panel B of Table 3 and plot the estimated return series from the CSV approach in Fig. 3. The correlation parameter ρ is positive for corn futures but it is negative for wheat futures. However, unlike the results of stock returns, none of the correlations are significant. Our CSV model is able to recover the censored returns and provide theoretically sound estimates of the true futures returns. Whereas limit hits are more spread out for stocks, they seem to cluster during a certain period for futures.

Panel C of Table 3 reports the estimation of the unconditional means, minimums, maximums, standard deviations, and excess kurtosis of log returns for the ostrich, discarding, CSV, and GMM approaches. Similar to our results for stocks, the CSV approach provides the lowest minimum and the highest maximum as censored returns are recovered. Our empirical evidence shows that the discarding approach has the lowest standard deviation because the most volatile movements are deleted. After considering the unobservability and spillover effect, our CSV approach gives higher estimation of standard deviation than does either the ostrich or the discarding approach. However, because limit hits for futures are not as frequent as those for stocks, the difference in standard deviation estimates for futures is not as noticeable as that for stocks. The CSV also provides the highest excess kurtosis that appropriately reveals the fat-tailed return distributions. Overall, our empirical results for futures are consistent with our simulation results.

6. Robustness check

To examine the robustness of our results, we compare the performance of our CSV approach with that of other existing models using the S&P 500 Index returns and the GARCH-generated returns. Unlike the TSMC, ASE, corn and wheat futures returns, the S&P 500 Index returns are not subject to price limits and thus true returns are observable. By artificially imposing price limits on the

Table 4
S&P500 index and GARCH-generated returns

Panel A

Returns	Bounds	# of limit hits	# of upper limit hits	# of lower limit hits
S&P500	±2%	5.72	3.20	2.52
GARCH	±2%	5.44	2.77	2.67

Panel B

Variable	S&P500			GARCH		
	Ostrich	Discarding	CSV	Ostrich	Discarding	CSV
μ	−9.374 (0.165)	−9.546 (0.161)	−9.314 (0.198)	−9.179 (0.191)	−9.323 (0.183)	−9.117 (0.224)
ϕ	0.696 (0.186)	0.690 (0.185)	0.737 (0.171)	0.754 (0.158)	0.736 (0.166)	0.787 (0.143)
σ_v	0.148 (0.030)	0.151 (0.031)	0.150 (0.032)	0.148 (0.029)	0.150 (0.030)	0.152 (0.031)
ρ	0.302 (0.194)	0.346 (0.202)	0.259 (0.175)	0.330 (0.196)	0.345 (0.200)	0.313 (0.186)

Panel C

Method	S&P500				GARCH			
	Unconditional mean		Unconditional standard deviation		Unconditional mean		Unconditional standard deviation	
	RMSE ($\times 10^{-4}$)	MPE (%)	RMSE ($\times 10^{-4}$)	MPE (%)	RMSE ($\times 10^{-4}$)	MPE (%)	RMSE ($\times 10^{-4}$)	MPE (%)
Ostrich	1.43	−16.3	17.29	−7.1	2.82	−4.0	19.14	−4.9
Discarding	4.37	−29.6	30.08	−14.3	7.63	−24.8	31.89	−11.5
CSV	2.32	−14.3	9.02	−1.2	2.90	1.4	13.76	1.4
GMM	1.43	−16.3	21.07	−8.5	2.82	−4.0	20.50	−5.5

This table reports results from the following stochastic volatility model using the S&P500 Index and the GARCH-generated returns.

$$r_t = e^{h_t/2}\varepsilon_t, \text{ and}$$

$$h_t = \mu + \phi(h_{t-1} - \mu) + \sigma_v v_t$$

where ε_t and $\sigma_v v_t$ are jointly normal with a mean vector $[0\ 0]'$ and the covariance matrix $\Sigma = \begin{pmatrix} 1 & \rho\sigma_v \\ \rho\sigma_v & \sigma_v^2 \end{pmatrix}$ for $t = 1,\ldots,n$, and $-1 < \phi < 1$. ρ is the correlation between ε_t and v_t. h_t is referred to as the latent log volatility that measures the amount of volatility on trading day t, and ϕ measures the autocorrelation present in the logged squared data. σ_v is considered the volatility of log-volatilities. $\text{RMSE} = \sqrt{\frac{1}{m}\sum_{i=1}^{m}\left(\hat{\theta}_i - \theta\right)^2}$, and $\text{MPE} = \frac{1}{m}\sum_{i=1}^{m}\left(\frac{\hat{\theta}_i - \theta}{\theta}\right)$, where $\hat{\theta}_i$ is a sample mean or standard deviation estimate of the ith sample and $m = 1000$ in the MCMC simulation, θ is the true sample mean or standard deviation. The CSV is our proposed censored stochastic volatility model, the Ostrich approach ignores price limits, and the discarding approach deletes limit-hit observations. Panel A reports the average number of limit hits for each 100 observations. Bounds are the price limit rates. Upper limit hits occur when the price hits the upper bound, and lower limit hits occur when the price hits the lower bond. Limit hits include both the upper and lower limit hits. Panel B reports the estimations of model parameters and Panel C reports the unconditional mean and standard deviation from our MCMC simulation. Posterior standard deviations are in parentheses. The generalized method of moments (GMM) estimator is that proposed by Wei and Chiang (2004).

S&P 500 Index returns, we are able to measure the errors of various models in estimating unconditional means and standard deviations as well as the model parameters. Besides, the use of GARCH-generated returns avoids potential bias in favor of our CSV approach, given that our earlier simulation data are generated from the SV model.

We obtain the daily S&P 500 Index from January 2, 1996 to March 31, 2006 from the Yahoo! Finance website. Given that returns may be serially correlated, we randomly draw a starting date and obtain a string of returns (100 observations) instead of drawing only one observation at a time. We repeat this process to obtain 1000 samples for our MCMC simulation. To be comparable to our earlier simulation, the price limits are set at 2% so that about 5% of the observations hit the price limits.

Table 4 reports the results for the S&P 500 Index and the GARCH-generated returns. Panel A shows that on average 5.72 (5.44) limit hits occur for every one hundred observations of the S&P 500 Index (GARCH-generated) returns. Similar to the results of stock returns, there are more upper limit hits than lower limit hits. Panel B reports the estimations of model parameters and Panel C reports the estimation of unconditional means and standard deviations of returns. For both the S&P 500 Index and GARCH-generated returns, our CSV outperforms other approaches in terms of the estimation of unconditional means and standard deviations, consistent with our previous results.

It should be noted that our CSV approach provides the estimation of unconditional means and standard deviations well even for GARCH-generated returns. Although it is still under debate whether the GARCH models or SV models are better in capturing the dynamics of financial time series, what we can assure is that our CSV approach performs well even if the return series follows the GARCH process. Overall, our results are robust.

7. Conclusion

In this paper, we propose a CSV approach to model the return process of assets that are subject to price limits. When price limits are reached, the observed prices are truncated and the equilibrium prices are unobservable, which makes it difficult to calculate the true returns and volatilities for any further financial analyses. Although several approaches and models have been suggested and applied, they are either unable to account for both the unobservability and the spillover effect associated with price limits or too computationally complicated to be feasible for simulation studies. Because we use a theoretically sound imputation method to generate fill-ins for censored observations, we are able to simulate the distribution of each unobserved return in a return series. These imputed returns and distributions are then appropriate for subsequent financial analyses, such as option pricing, risk management, and asset allocations, as we have demonstrated in calculating the preceding unconditional standard deviations.

To assess the performance of our model, we compare it with the ostrich, discarding, and GMM approaches through an MCMC simulation. The results of our simulation analysis suggest that our model outperforms other approaches with respect to the estimates of model parameters, unconditional means, and standard deviations. We further apply our algorithm to model the returns of two stocks and two futures contracts that are subject to price limits and find that the results are consistent with our simulation outcomes. Overall, our model is not only theoretically sound but also empirically applicable.

Because of the better performance of our model, we suggest it should be used to model return series that are subject to price limits. Many financial markets in the world impose price limits; therefore, international portfolio managers might use our model to obtain better estimates of risk for financial assets. Ignoring the existence of price limits or deleting limit-hit observations leads to an underestimation of asset risks, which may incur unnecessary losses due to mismanaged portfolios and asset allocations. Corporate managers also might use our model to create a full return series for their company stocks to obtain appropriate estimates of the cost of capital for their capital budgeting analyses. To study the impact of price limit rules, regulators and researchers could use our model to conduct extensive simulation studies. Although it is relatively difficult to test the effectiveness of price limits, because we simply do not know what would have happened without price limits, the estimation from our model at least provides researchers with a benchmark that can be used to evaluate the performance of price limits.

It should be noted that our CSV model has its limitation. For example, in a market where price limits may be widen on a given day, our model is not appropriate. A good example is the Spanish Stock Exchange where price limits trigger trading halts. Regulators constantly monitor the market and decide whether to widen the price limits after trading halts. Our model treats price limits as given on any given day and thus is unable to capture the widened price limits when that happens.

Appendix A. Implementation of Steps A and B

A.1. Implementation of Step A — ε_t and v_t are independent

The joint prior density of the set θ, with the assumption of prior independence, is

$$p(\theta) = p(\mu)\cdot p(\sigma_v^2)\cdot p(\phi)\cdot p(h_0|\mu,\sigma_v^2)\cdot \prod_{t=1}^{n} p(h_t|h_{t-1},\mu,\sigma_v^2,\phi). \tag{A1}$$

Let $R=\{r_t^o|t\le T_o\}\cup R^*$ be the set of all equilibrium returns. From Eq. (2) and conditional independence, the "latent" likelihood function is $p(R|\theta)=p(r_1,r_2,\dots,r_n|\theta)=\prod_{t=1}^{n} f(r_t|\theta)$ where $f(r_t|\theta)$ is a normal density with a mean 0.0 and variance e^{h_t}. Hence, the posterior of θ given R is $p(\theta|R)\propto p(R|\theta)\cdot p(\theta)$. Step A is broken down into (A-1) sample from $f(\sigma_v^2|\mu,\phi,h_t,R)$, (A-2) sample from $f(\mu|\phi,\sigma_v^2,h_t,R)$, (A-3) sample from $f(\phi|\mu,\sigma_v^2,h_t,R)$, and (A-4) sample from $f(h_t,t=1,2,\dots,n|\mu,\sigma_v^2,\phi,R)$. For Steps A-1 and A-2, we employ the standard natural conjugate priors as in Kim et al. (1998), and the resulting posteriors are of standard form. A beta prior for $\phi^*=0.5(\phi+1)$ is used in Step A-3 and its posterior is proportional to $\Psi(\phi)\cdot N(\hat{\phi},\hat{V})$, where $\Psi(\phi)$ is uniformly bounded, and $N(\hat{\phi},\hat{V})$ is a normal density, see Chib and Greenberg (1994).

The accept–reject procedure (Ripley, 1987) is used in Step A-4 for $h_t, t=1,2,\ldots,n$. To sample from the joint posterior distribution $f(h_t, t=1,2,\ldots,n|\mu,\sigma_v^2,\phi,R)$, we sample each h_t from $f(h_t|h_{\backslash t},\mu,\sigma_v^2,\phi,R)$, where $h_{\backslash t}=\{h_1,h_2,\ldots,h_n\}\backslash\{h_t\}$. This operation is carried out n times for each sweep to obtain a random sample from the joint posterior distribution. Note that $f(h_t|h_{\backslash t},\mu,\sigma_v^2,\phi,R) \propto f(h_t|h_{\backslash t},\mu,\sigma_v^2,\phi) f(r_t|h_t,\mu,\sigma_v^2,\phi)$. In turn, we show that $f(h_t|h_{\backslash t},\mu,\sigma_v^2,\phi)=f(h_t|h_{t-1},h_{t+1},\mu,\sigma_v^2,\phi)=N(\mu_h,\sigma_h^2)$ and $f(r_t|h_t,\mu,\sigma_v^2,\phi)$ is dominated by a function $g(\cdot)$, where $\mu_h=\frac{\phi(h_{t+1}+h_{t-1})+\mu(1-\phi)^2}{1+\phi^2}$, $\sigma_h^2=\frac{\sigma_v^2}{1+\phi^2}$, and $\log g=-\frac{h_t}{2}-\frac{r_t^2}{2}\left[e^{-\mu_h}(1+\mu_h)-h_t e^{-\mu_h}\right]$. We sample a value h_t from $N(\mu_h,\sigma_h^2)$ and accept it with probability $f(r_t|h_t,\mu,\sigma_v^2,\phi)/g$. If the value is rejected, we redraw a value from the normal density $N(\mu_h,\sigma_h^2)$.

A.2. Implementation of Step A — ε_t and v_t are correlated

The sampling processes described above for Steps A-2, A-3, and A-4 are applicable here. For Step A-1, Jacquier et al. (2004) suggest a transform of Σ so that conjugate priors can be used on the elements of Σ.

A.3. Implementation of Step B

Let r_k be a latent return in R^*. Given θ, R^o, and all latent returns in R^* except r_k (i.e., $R^*-\{r_k\}$), the conditional density of r_k is a normal density truncated from below if the price hits the upper limit at either time k and/or time $k-1$, but it is a normal density truncated from above if the price hits the lower limit at time k and/or $k-1$. Specifically,

$$p\left(r_k|\theta,R^o,R^*-\{r_k\}\right)=\begin{cases} f_u(r_k)=f_N(r_k)I[r_k^o-E_{k-1},\infty), & \text{if } k\in T_u \\ f_d(r_k)=f_N(r_k)I\left(-\infty,r_k^o-E_{k-1}\right], & \text{if } k\in T_d, \end{cases} \tag{A2}$$

where f_N is the normal density with a mean 0.0 and variance e^{h_t}, and $I_A=1$ if $r_k\in A$ and 0.0 otherwise. To generate the latent returns at the ith iteration, we sequentially sample r_{j_k} from Eq. (A2), where $r_{j_k}\in R^*$, $k=1,2,\ldots,c$, and c is the total number of limit hits. By cycling through all $r_{j_k}\in R^*$, we obtain a sample of $R^*=\{r_{j_1},r_{j_2},\ldots,r_{j_c}\}$ from $p(R^*|R^o,\theta)$. By default, we always choose to start a time series with a nonlimited observation, so $E_1=0$.

Appendix B. Adjustments of price limits on the Taiwan Stock Exchange

Date of adjustment	Price limits (based on previous closing price)
02/09/1962	5% upward and downward
04/09/1973	3% upward and downward
08/07/1973	5% upward and downward
02/19/1974	5% upward and 1% downward
03/09/1974	5% upward and one tick size downward
04/15/1974	1% upward and downward
05/21/1974	3% upward and downward
06/17/1974	5% upward and downward
12/19/1978	2.5% upward and downward
01/05/1979	5% upward and downward
10/27/1987	3% upward and downward
11/14/1988	5% upward and downward
10/11/1989	7% upward and downward
09/27/1999	7% upward and 3.5% downward
10/08/1999	7% upward and downward

Appendix C. Price limits of futures contracts on the Chicago Board of Trade

Start date	End date	Price limits per bushel	Price limits per contract
Panel A: corn futures			
11/16/93		$0.12	$600
06/01/94	06/02/94	$0.18	$900
06/03/94		$0.12	$600
06/21/94	06/22/94	$0.18	$900
06/23/94		$0.12	$600
04/30/96	05/01/96	$0.18	$900
05/02/96		$0.12	$600
05/21/96	05/22/96	$0.18	$900
05/23/96		$0.12	$600

(continued on next page)

Appendix C (*continued*)

Start date	End date	Price limits per bushel	Price limits per contract
Panel A: corn futures			
05/30/96	05/31/96	$0.18	$900
06/03/96		$0.12	$600
06/04/96	06/05/96	$0.18	$900
06/06/96		$0.12	$600
06/11/96	06/12/96	$0.18	$900
06/13/96		$0.12	$600
06/19/96	06/20/96	$0.18	$900
06/21/96		$0.12	$600
07/01/96	07/02/96	$0.18	$900
07/03/96		$0.12	$600
07/17/96	07/18/96	$0.18	$900
07/19/96		$0.12	$600
07/23/96	07/24/96	$0.18	$900
07/25/96		$0.12	$600
08/13/96	08/14/96	$0.18	$900
08/15/96	08/12/97	$0.12	$600
Panel B: wheat futures			
02/17/92		$0.20	$1,000
04/11/96	04/15/96	$0.30	$1,500
04/16/96		$0.20	$1,000
04/17/96	04/19/96	$0.30	$1,500
04/22/96		$0.20	$1,000
04/23/96	04/25/96	$0.30	$1,500
04/26/96		$0.20	$1,000
04/30/96	05/02/96	$0.30	$1,500
05/03/96		$0.20	$1,000
05/09/96	05/13/96	$0.30	$1,500
05/14/96		$0.20	$1,000
05/30/96	06/03/96	$0.30	$1,500
06/04/96	04/14/97	$0.20	$1,000

References

Andersen, T.G., Chung, H.J., Sorensen, B.E., 1999. Efficient method of moments estimation of a stochastic volatility model: a Monte Carlo study. Journal of Econometrics 91, 61–87.

Andersen, T.G., Bollerslev, T., Diebold, F.X., Ebens, H., 2001. The distribution of realized stock return volatility. Journal of Financial Economics 61, 43–76.

Berkman, H., Steenbeek, O.W., 1998. The influence of daily price limits on trading in Nikkei futures. Journal of Futures Markets 18, 265–279.

Calzolari, G., Fiorentini, G., 1998. A Tobit model with GARCH errors. Econometric Reviews 17, 85–104.

Chen, H., 1998. Price limits, overreaction, and price resolution in futures markets. Journal of Futures Markets 18, 243–263.

Chiang, R., Wei, K.C.J., Wu, S., 1990. Price limits in Taiwan and risk-return estimation. In: Rhee, S.G., Chang, R.P. (Eds.), Pacific-Basin Capital Markets Research. Elsevier Science Publishers BV, North-Holland.

Chib, S., Greenberg, E., 1994. Bayes inference in regression models with ARMA (p, q) errors. Journal of Econometrics 64, 183–206.

Cho, D.D., Russell, J., Tiao, G.C., Tsay, R., 2003. The magnet effect of price limits: evidence from high-frequency data on Taiwan Stock Exchange. Journal of Empirical Finance 10, 133–168.

Chou, P.H., 1997. A Gibbs sampling approach to the estimation of linear regression models under daily price limits. Pacific-Basin Finance Journal 5, 39–62.

Chung, J., Gan, L., 2005. Estimating the effect of price limits on limit-hitting days. Econometrics Journal 8, 79–96.

Durham, G.B., 2007. SV mixture models with application to S&P 500 index returns. Journal of Financial Economics 85, 822–856.

George, T.J., Hwang, C.Y., 1995. Transitory price changes and price-limit rules: evidence from the Tokyo Stock Exchange. Journal of Financial and Quantitative Analysis 30, 313–327.

Gilks, W.R., Wild, P., 1992. Adaptive rejection sampling for Gibbs sampling. Journal of the Royal Statistical Society — Series C 41, 337–348.

Gilks, W.R., Roberts, G.O., Sahu, S.K., 1998. Adaptive Markov chain Monte Carlo through regeneration. Journal of the American Statistical Association 93, 1045–1054.

Jacquier, E., Polson, N.G., Rossi, P.E., 1994. Bayesian analysis of stochastic volatility models. Journal of Business and Economic Statistics 12, 371–389.

Jacquier, E., Polson, N.G., Rossi, P.E., 2004. Bayesian analysis of stochastic volatility models with fat-tails and correlated errors. Journal of Econometrics 122, 185–212.

Kim, K.A., Rhee, S.G., 1997. Price limit performance: evidence from the Tokyo Stock Exchange. Journal of Finance 52, 885–901.

Kim, Y.H., Yang, J.J., 2004. What makes circuit breakers attractive to financial markets? A survey. Financial Markets, Institutions, and Instruments 13, 109–146.

Kim, S., Shephard, N., Chib, S., 1998. Stochastic volatility: likelihood inference and comparison with ARCH models. Review of Economic Studies 65, 361–393.

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.

Kon, S.J., 1984. Models of stock returns — a comparison. Journal of Finance 39, 147–165.

Lee, L.F., 1999. Estimation of dynamic and ARCH Tobit models. Journal of Econometrics 92, 355–390.

Lee, S.B., Kim, D.J., 1997. Price limits and beta. Review of Quantitative Finance and Accounting 9, 35–52.

Ma, C.K., Rao, R.P., Sears, R.S., 1989. Volatility, price resolution, and the effectiveness of price limits. Journal of Financial Services Research 3, 165–199.

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

Meyer, R., Yu, J., 2000. BUGS for a Bayesian analysis of stochastic volatility models. Econometrics Journal 3, 198–215.

Miller, M.H., 1989. Commentary: volatility, price resolution, and the effectiveness of price limits. Journal of Financial Services Research 3, 201–203.

Morgan, I.G., Trevor, R.G., 1999. Limit moves as censored observations of equilibrium futures price in GARCH processes. Journal of Business and Economic Statistics 17, 397–408.

Nardari, F., Scruggs, J.T., 2007. Bayesian analysis of linear factor models with latent factors, multivariate stochastic volatility, and APT pricing restrictions. Journal of Financial and Quantitative Analysis 42, 857–891.
Neal, R.M., 1997. Markov chain Monte Carlo methods based on 'slicing' the density function. Technical Report No. 9722. Department of Statistics, University of Toronto.
Ripley, B.D., 1987. Stochastic Simulation. Wiley, New York.
Ritter, C., Tanner, M.A., 1992. Facilitating the Gibbs sampler: the Gibbs stopper and the Griddy–Gibbs sampler. Journal of the American Statistical Association 87, 861–868.
Roll, R., 1984. Orange juice and weather. American Economic Review 74, 861–880.
Subrahmanyam, A., 1994. Circuit breakers and market volatility: a theoretical perspective. Journal of Finance 49, 237–254.
Sutrick, K.H., 1993. Reducing the bias in empirical studies due to limit moves. Journal of Futures Markets 13, 527–543.
Tanner, M.A., 1993. Tools for Statistical Inference: Methods for the Exploration of Posterior Distributions and Likelihood Functions, 2nd ed. Springer-Verlag.
Tanner, M.A., Wong, W.H., 1987. The calculation of posterior distributions by data augmentation. Journal of the American Statistical Association 82, 528–540.
Tauchen, G.E., Pitts, M., 1983. The price variability–volume relationship on speculative markets. Econometrica 51, 485–505.
Taylor, S.J., 1982. Financial returns modeled by the product of two stochastic processes—a study of the daily sugar prices 1961–75. In: Anderson, O.D. (Ed.), Time Series Analysis: Theory and Practice. North-Holland, Amsterdam.
Wei, S.X., 1999. A Bayesian approach to dynamic Tobit models. Econometric Reviews 18, 417–439.
Wei, S.X., 2002. A censored-GARCH model of asset returns with price limits. Journal of Empirical Finance 9, 197–223.
Wei, K.C.J., Chiang, R., 2004. A GMM approach for estimation of volatility and regression models when daily prices are subject to price limits. Pacific-Basin Finance Journal 12, 445–461.