Risk Management during the Real Estate Bubble: GARCH or Stable Distributions?

David Basterfield and Thomas Bundt

Abstract

This paper backtests Value-at-Risk (VaR) for the Stable Paretian and GARCH models applied to the real estate bubble of 2005-2008. Specifically, we use a rolling time-varying quantile estimation method to backtest Value-at-Risk (VaR) on a widely-held real-estate ETF. Our statistical analysis allows us to test for both distributional assumptions and a model's ability to track volatility clustering. We find that neither the Stable Paretian nor GARCH model performs satisfactorily for both 95% and 99% VaR and over both crisis and pre-crisis periods. In some cases our rolling time-varying parameter estimation methodology allows the Stable model to successfully track volatility clustering, a procedure simpler to apply than standard GARCH models. Our results are sensitive to the length of the conditioning window, with both models doing better for the 50-day window relative to 100 and 200-day windows. Finally, of particular interest is the Stable model's high frequency of rejecting the unconditional coverage null, suggesting the Stable distribution poorly fit the data.

I. Purpose and Overview

This paper backtests Value-at-Risk (VaR) for the Stable Paretian and GARCH models applied to the real estate bubble of 2005-2008. Specifically, we use a time-varying quantile estimation method to backtest VaR on a widely-held real-estate exchange-traded fund (ETF). Our rolling estimation procedure allows the Stable model to track volatility clustering facilitating a direct comparison with the industry standard GARCH model. Our statistical analysis allows us to test for both distributional assumptions and a model's ability to track volatility clustering.

There is a growing literature on risk management applications of Stable distributions (Mittnik and Rachev 2000). (Young and Graff, 1995) show real estate return distributions follow a Stable distribution relative to the normal. (Basterfield and Bundt 2009) show the Stable model typically performs better during periods of high volatility at extreme quantiles relative to the GARCH when applied to exchange rate data during the Asian currency crisis. Yet, the application of stable distributions in finance is not robust and their usefulness, to some, has been discounted. For instance, conventional wisdom argues that while Stable distributions typically outperform the normal distribution, they tend to overestimate tail thickness in finance applications (Mittnik and Rachev 1993).

We take the position that the debate on the risk-management applications of the Stable distributions is unresolved. Accordingly, we extend the current literature in two ways: (1) we examine the efficacy of Stable distributions during the recent real estate bubble, and (2) we use a rolling time-varying parameter estimation methodology allowing us to compare Stable and GARCH models directly using conditional coverage failure rate analysis. Our goal is to contrast and compare Stable with GARCH models during a period of time of a widely-recognized

David Basterfield, Associate Professor of Finance, Department of Economics and Business Administration, Hillsdale College, 33 East College, Hillsdale MI, 49242, (517) 607-2412, dbasterfield@hillsdale.edu
Thomas Bundt, Associate Professor of Finance and Quantitative Analysis, Department of Economics and Business Administration, Hillsdale College, 33 East College, Hillsdale MI, 49242, (517) 607-2419, tbundt@hillsdale.edu

financial *bubble*. Accordingly, our research is aimed at testing VaR models when they matter the most, during a financial crisis.

II. Data and Descriptive Statistics

Daily trading data on real-estate investment trusts (REITs) is available in the form of ETFs. Based on having the highest trading volume, we selected the iShares Dow Jones US Real Estate ETF, trading symbol IYR on NYSEArca, which tracks the Dow Jones US Real Estate Index. We collected daily trading data on IYR over the ten-year period 6/19/2000-12/15/2009 available from Yahoo finance. The daily close is plotted in Figure 1. The nature of the real estate bubble is clearly shown with index values doubling over the 2004-2006 period, only to reverse during 2007-2008.

To determine the appropriate time period to split our sample into pre-crisis and crisis periods, we plotted a 30-day moving average of volatility, calculated as daily standard deviation, shown in Figure 2. Clearly volatility begins an upward spiral around 8/1/2007. Accordingly, we split our full sample into (1) a pre-crisis period 6/19/2000-7/31/2007, and (2) a crisis period 8/1/2007-12/15/2009. This will allow us to study a low- and high-volatility subsample allowing relative performance comparisons between the two VaR models.

We converted daily data into a continuously compounded return series. Table I compares summary statistics over the two periods. As reported in Table I, both periods are characterized by non-normal distributions as reported by the Jarque-Bera test. In addition, the change in volatility over the two periods is quite dramatic, with a 400% increase in volatility into the crisis sub-period.

III. Stable Index Estimates

(Lévy, 1924) showed tails of non-Gaussian stable distributions asymptotically follow the law of Pareto, giving rise to the term "Stable Paretian." Stable distributions have four parameters: characteristic exponent or index of stability $\alpha \in (0, 2]$, skewness parameter β, scale parameter γ, and location parameter δ. The normal characteristic function requires restrictions

$\alpha = 2$, $\beta = 0$, $\gamma = \sigma^2/2$, and $\delta = \mu$. The key parameter governing the rate of decline of tail mass is the index of stability α. An important property of stable laws is moments of order $r > \alpha$ do not exist. Accordingly, if $\alpha < 1$, tails are of such large mass that expected value is undefined, i.e., moment-generating integrals are unbounded. For $\alpha < 2$, variance and covariance do not exist.

Table II reports 95% confidence intervals for stable index estimates over both periods using maximum likelihood[21]. Since the upper limits for both periods are less than 2, we reject the null of a finite variance process in both periods. None of our lower bounds are below 1, so we cannot reject the null of a finite mean. Note the 95% confidence interval for the pre-crisis period is much broader than in the crisis period revealing our alpha estimates during the crisis period are relatively imprecise. Nevertheless, the mean index estimate is smaller in the crisis sub-period reflecting the increased volatility characteristic of the crisis or high-volatility sub-period.

---

21 We estimated Stable indexes using STABLE available at http://academic2.american.edu/~jpnolan/. A general description of the software is found in Nolan (1997).

IV. VaR Models

Market risk analysis in finance typically involves mappings from extreme quantiles of return distributions to exposure measures such as Value at Risk. Specifically, VaR_p is the *p*-percentage quantile of our returns distribution, exceeded with probability *(1 – p)*, conditional on information in the previous period Ω_{t-1}:

$$(1) \qquad VaR_p = F^{-1}_{t|\Omega_{t-1}}(p)$$

Accordingly, VaR prediction amounts to accurately forecasting returns distribution quantiles.

We backtested two VaR models: GARCH (1, 1), and Stable Paretian. The GARCH (1,1) model uses a conditional variance process to obtain daily volatility estimates:

$$(2) \qquad \sigma^2_{t|\Omega_{t-1}} = \delta_o + \delta_1 \varepsilon_t^2 + \delta_2 \sigma^2_{t-1}$$

Accordingly, the conditional standard deviation for day t, σ_t, was estimated on a rolling daily basis using MLE and equation (2). Quantile estimates were then obtained using:

$$(3) \qquad q_{P,t} = \Phi_P^{-1}\left(\sigma_{t|\Omega_{t-1}}\right)$$

where $q_{P,t}$ is the *p*-percentage quantile forecast for time period *t*, Φ_P^{-1} is the standard normal inverse function for probability *p*, and $\sigma_{t|\Omega_{t-1}}$ is the volatility forecast conditioned on information in period *t-1*.

Rolling quantile estimates for stable distributions were calculated directly using STABLE:

$$(4) \qquad q_{p,t} = S\left(\alpha, \beta, \gamma, \delta; \Omega_{t-1}\right).$$

V. Backtesting Procedures and Results

We backtested our VaR models using daily performance comparisons used by trading firms and mandated by regulators.[22] These procedures amount to failure rate analysis applied to VaR exceptions, defined as events when actual VaR exceeds its forecast. We controlled for the clustering of VaR exceptions by following conditional coverage tests of (Christoffersen 1998), which test whether a sequence of VaR exceptions constitute an i.i.d. Bernoulli sequence. Specifically, conditional coverage tests examine the joint null of the unconditional coverage rate

22 For a review of mandated backtesting procedures, see (Jorion 2007) pages 142-153. Note that our tests and model rankings are based on 50- and 200-day windows. We also tested a 100-day window and found it differed only slightly from the 200-day window.

and independent exceptions to VaR. The conditional coverage null is $E[I_t|\Omega_{t-1}] = p$, where p is the coverage rate, I_t indicates an exception to VaR at time t, and Ω_{t-1} is the information set. Specifically, I_t, the exceptions indicator variable for day t, is set to one when the actual loss exceeds the forecast VaR for that day, and zero otherwise. Following (Kupiec 1995), the unconditional null is $E[I_t] = p$ which ignores how exceptions cluster in time. The null of independent exceptions is tested against a first-order Markov process. Accordingly, rejection of the conditional coverage null can be explained by 1) an inappropriate distributional assumption, or 2) clustering of exceptions, or 3) both.

Like the normal, the Stable distribution is typically applied under the i.i.d. assumption and therefore offers no dynamic structure to the researcher to model volatility clustering. To remedy this we employed a rolling time-varying parameter estimation technique to condition daily Stable VaR forecasts. Specifically, we used 50, 100, and 200-day windows to condition our quantile estimates. This rolling estimation approach allows us to compare models not designed to specifically model heteroscedasticity, notably the Stable Paretian model, with the GARCH benchmark, which is specifically designed to track volatility clustering[23]. In addition, our approach requires less computational effort and parameters estimates than the industry standard GARCH model.

Table III presents likelihood ratio values for the conditional coverage null for each model at the 95% and 99% VaR using a 50-day conditioning window over the pre-crisis and crisis periods. In the pre-crisis period, we reject 95% and 99% conditional coverage nulls for both models. Additional analysis concluded the stable model conditional coverage null was rejected for 99% VaR because the unconditional coverage null was rejected, a surprise given our prior that Stable models perform best at extreme quantiles. In addition, the fact that the GARCH model rejects the unconditional coverage null at the 95% quantile during a period of relative calm is surprising since our priors were the GARCH model would outperform the Stable at the 95% quantile. Finally, the failure of GARCH to track volatility clustering for both 95% and 99% quantiles, for which the model is designed, was a surprise.

Over the crisis period we fail to reject the conditional coverage null for both models at 95% and 99% quantiles. Comparing likelihood ratio values, which follow a chi-square distribution, we can rank models where the lower the likelihood ratio value the better the model fit. Based on likelihood ratios, further analysis revealed the GARCH model outperforms the Stable for 95% whereas it underperforms the Stable for 99% VaR confirming our prior that Stable models perform best at extreme quantiles. It is interesting that our time-varying estimation approach tracks volatility clustering for the Stable model during the crisis period, whereas it fails during the pre-crisis period.

Table IV presents statistical tests of coverage rates using a 200-day conditioning window.
For the pre-crisis period, we reject the conditional coverage null for the GARCH model at both quantiles. Further analysis reveals both the unconditional coverage and independence nulls were rejected for both 95% and 99% quantiles for the GARCH model, a surprise given the GARCH

23 Our rolling parameter estimation approach is consistent with time-varying parameters and is similar to a Bayesian updating model. For example, (Berkowitz and O'Brien 2002) update ARMA-GARCH model parameters on a daily basis for internal bank VaR data.

model is designed to track volatility clustering. On the other hand, for the pre-crisis period, we reject the conditional coverage null for the Stable model at the 95% quantile, but fail to reject for the 99% quantile. This result confirms our priors which expected the Stable model to outperform the GARCH at extreme quantiles.

For the crisis period, the GARCH model fails to reject the conditional coverage null for both quantiles, whereas we reject the conditional coverage null for both quantiles for the Stable model. Analysis revealed, for both quantiles, the Stable model rejected the unconditional null indicating it was an inappropriate distribution but failed to reject the independence null. This is a surprising result since the stable model typically performs better during periods of high volatility. Finally, the superiority of the GARCH model in the crisis period was also unexpected. Accordingly, for the 200-day conditioning window, the significant result is that the GARCH model outperforms the stable at extreme quantiles during a period of relatively high volatility.

VI. Conclusion

This paper evaluates the predictive quality of day-ahead VaR forecasts using a time-varying parameter estimation approach applied to daily real-estate index data during the real-estate bubble. Specifically, we backtest daily-ahead VaR forecasts for two models: Stable Paretian and GARCH. Our testing procedures allow separate evaluation of distributional assumptions from model dynamics associated with volatility clustering.

Testing Stable and GARCH models during periods of financial volatility is important since, at its heart, stable distributions are capable of modeling rare events. Indeed, stable models can forecast infinite-variance return distributions. On the other hand, GARCH models successfully track volatility clustering and are widely employed on Wall Street despite their adherence to smooth continuous well-defined return distributions. Examining how each model performs during financial bubbles is important in future risk management applications.

What insights stem from our research? First, our rolling estimation approach allows the Stable model to successfully track volatility clustering, a procedure simpler to apply than standard GARCH models. Second, our results are sensitive to the length of the conditioning window, with both models doing better for the 50-day window relative to 100 and 200-day windows. Third, neither model performs well over both periods and over quantiles, with the GARCH marginally outperforming the Stable. However, for the 50-day conditioning window during the crisis period, our results confirm what was expected, the stable model outperforms the

GARCH for extreme quantiles although neither model is satisfactory for the pre-crisis period. Of particular interest is the Stable model’s high frequency of rejecting the unconditional coverage null, suggesting the Stable distribution poorly fit the data. What can explain the relatively poor performance of the Stable model? We suspect that the alpha values indicate tail masses are not so large as to provide a distribution advantage to the Stable model. Accordingly, areas of future research include re-examining the role of the normal model in the real-estate bubble.

What do our results imply about the use of VaR models in a financial bubble? The poor performance of both Stable and GARCH models is likely to disappoint adherents to VaR and Stable models as they don't perform as expected, i.e., Stable models perform relatively poorly at extreme quantiles, where they are expected to have a comparative advantage relative to GARCH models. Accordingly, additional research on VaR and Stable models during financial bubbles is required to resolve some of these issues.

Figures and Tables

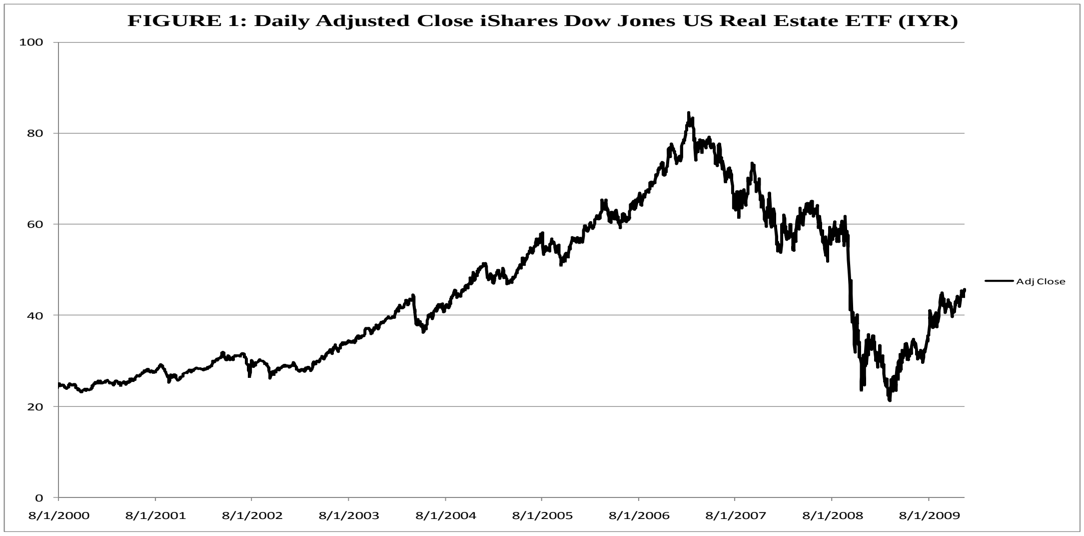

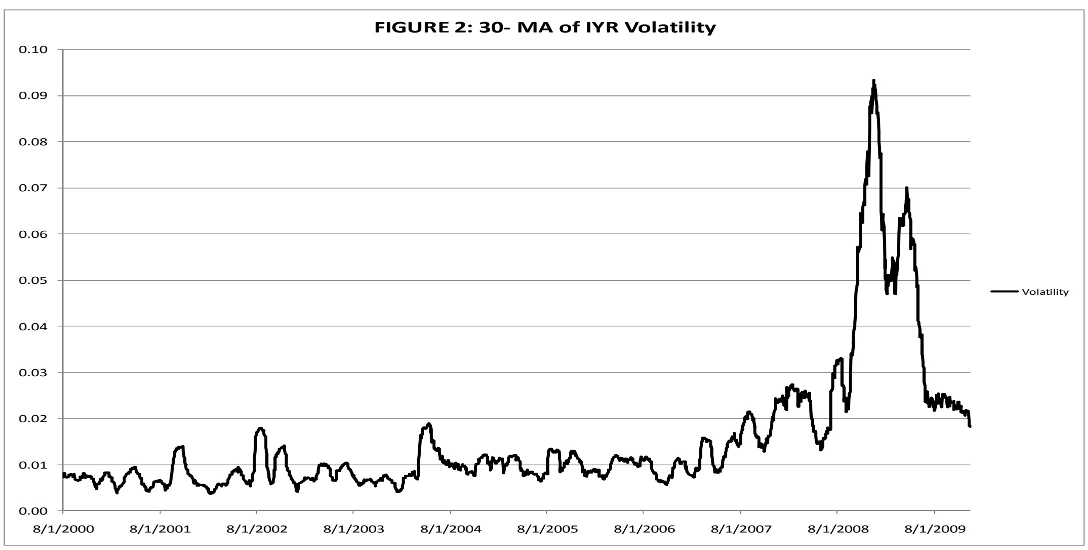

TABLE I: Summary Statistics		
Descriptive Statistic	**Pre-crisis Period 6/19/2000-7/31/2007**	**Crisis Period 8/1/2007-12/15/2009**
Mean	0.06%	-0.06%
Standard Deviation	0.96%	3.93%
Skewness	-0.46	-0.26
Kurtosis	5.02	6.64
Range	0.09	0.38
Jarque-Bera p-value	0.00	0.00

TABLE II: Stable Index Estimates	
Precrisis Period 6/19/2000-7/31/2007	**Crisis Period 8/1/2007-12/15/2009**
[1.7968, 1.9098]	[1.4146, 1.688]

TABLE III: Conditional Coverage Statistics 50-Day Window		
Quantile	**Pre-crisis Period 6/19/2000-7/31/2007**	**Crisis Period 8/1/2007-12/15/2009**
95%		
GARCH	25.7766*	1.5575
Stable	28.2409*	2.0515
99%		
GARCH	36.6570*	3.7881
Stable	10.2120*	2.5825
* denotes rejection of conditional coverage null at 5% significance level		

TABLE IV: Conditional Coverage Statistics 200-Day Window		
Quantile	**Pre-crisis Period 6/19/2000-7/31/2007**	**Crisis Period 8/1/2007-12/15/2009**
95%		
GARCH	15.5282*	3.7450
Stable	35.2663*	10.9987*
99%		
GARCH	22.4314*	5.1862
Stable	6.0848	10.3951*
* denotes rejection of conditional coverage null at 5% significance level		

References

Basterfield, D. and T. Bundt, 2009, Backtesting Value at Risk during the Asian Currency Crisis: Stable Paretian versus GARCH Models, *Journal of Global Business Development*, 1 (1), 125-137.

Berkowitz, J. and J. O'Brien, 2002, How Accurate are Value-At-Risk Models at Commercial Banks? *The Journal of Finance*, 57, 1093-1111.

Christoffersen, P., 1998, Evaluating Interval Forecasts, *International Economic Review*, 39, 841-862.

Jorion, P., 2007, *Value at Risk*, 3rd Ed., McGraw-Hill.

Kupiec, P., 1995, Techniques for Verifying the Accuracy of Risk Measurement Models, *The Journal of Derivatives*, 2, 73-84.

Lévy, P., 1924, Theorie Des Erreurs La Loi De Gauss Et Les Lois Exceptionelles, *Bulletin de la Societe de France*, 52, 49-85.

Mittnik, S. and S. Rachev, 1993, Modeling Asset Returns with Alternative Stable Distributions, *Econometric Reviews*, 12, 261-330.

Mittnik, S. and S. Rachev, 2000, *Stable Paretian Models in Finance*, Wiley

Nolan, J., 1997, Numerical Computation of Stable Densities and Distribution Functions, *Communications in Statistics – Stochastic Models*, 13(4), 759-774.

Young, M. and R. Graff, 1995, Real Estate is Not Normal: A Fresh Look at Real Estate Return Distributions, *Journal of Real Estate Finance and Economics*, 10, 225-259