The Role of Restrictions in Estimating the Risk-Return Tradeoff

Myung D. Park
Department of Agricultural and Applied Economics
The University of Georgia
Athens, GA 30602-7509
email: mdpark@uga.edu

and

Jeffrey H. Dorfman
Department of Agricultural and Applied Economics
The University of Georgia
Athens, GA 30602-7509
email: jdorfman@uga.edu

August 13, 2008

Abstract

This paper investigates the risk-return tradeoff using a GARCH-M model. Previous studies on stock returns have found mixed results. Bayesian inference is employed to avoid the difficulty of inequality restrictions in classical inference. Our method allows restrictions directly on volatility, rather than on a set of model parameters. The sign of the risk-return tradeoff follows the mainstream of finance theory when the effect of unnecessary restrictions is removed.

Keywords: Bayesian econometrics; GARCH-M; Risk-return tradeoff; TARCH

JEL classification: C11; G10

1. Introduction

Since Autoregressive Conditional Heteroskedasticity in mean (ARCH-M) and GARCH in Mean (GARCH-M) models were introduced by Engle, Lilien, and Robins (1987) and Bollerslev, Engle, and Wooldridge (1988), a number of authors have explored the risk return tradeoff using a GARCH-M framework. Following ARCH and GARCH specifications of variance equations, other conditional volatility specifications in the GARCH-M framework have been created to capture the asymmetry which is often called the leverage effect of U.S. stock returns. For instance, standard GARCH-M is employed by French, Schwert, and Stambaugh (1987). Campbell and Hentschel (1992) and Sentana (1995) use QGARCH to find the relation between return and its conditional volatility. Baillie and DeGennaro (1990) consider the federal fund rate a factor of conditional variance and show the positive risk return tradeoff. TARCH and EGARCH are also employed by Nelson (1991) and Glosten, Jagannathan, and Runkle (1993).

Some researchers have found the expected positive risk-return tradeoff. However, some have found a negative mean variance tradeoff which is contrary to the mainstream of finance theory. The sign of the relationship appears sensitive to the volatility specification. This negative relation of risk and return is usually found with asymmetric specifications of conditional volatility like Exponential GARCH (EGARCH : Nelson, 1991) and Threshold ARCH model (TARCH : Rabemananjara and Zakoian, 1995; Glosten, Jagannathan, and Runkle, 1993). Lundblad (2007) reviews more results and suggests small sample size is the culprit behind negative and statistically weak results.

We propose a different explanation. Inequality restrictions are commonly imposed on the coefficients in the estimation of conditional volatility to ensure positive variance estimates. The estimate of the risk-return tradeoff might suffer from the restricted parameters. As Geweke's ARCH study on Bayesian inference (1989) points out, classical inference subject to inequality restrictions is an ill-defined and unsolved problem. We therefore turn to practical solutions employing explicit Bayesian methods and a restricted prior. Engle and Ng (1993) compare the performances of popular conditional volatility models and suggest that the TARCH specification is the best for evolving conditional volatility. Following their suggestion, we employ the TARCH model here.

The results we find suggest the evidence of the risk return tradeoff follows what finance theory expects. We further find the results to be sensitive to the nature of restrictions imposed to ensure a statistically plausible model. Employing restrictions directly on volatility instead of on a set of GARCH model parameters appears to solve the problem of empirically suspicious results.

2. Data description and Econometric specification.

In this paper, we use U.S. equity market monthly returns from the Center for Research in Security Prices (CRSP). Excess equity return is the difference between total market return and risk free asset return. Monthly returns are the value-weighted CRSP index of NYSE, AMEX, and Nasdaq. One-month Treasury bill return is used as the risk-free asset return. The excess return is the monthly value-weighted CRSP index minus the one-month Treasury bill return for the period 1926 to 2007 ($T = 983$).

The GARCH-M specification consists of two parts: the mean equation and the variance equation. The mean equation is

$$r_t = \mu + \lambda h_t + \epsilon_t, \qquad \epsilon_t \sim N(0, h_t) \tag{2.1}$$

where r_t is the monthly excess return, μ is the risk free return, h_t is the conditional variance and λ is the coefficient that represents the risk-return tradeoff.

The variance equation follows the TARCH specification which allows different reactions to volatility depending on the sign of past errors,

$$h_t = \omega + \alpha \epsilon_{t_1}^2 + \gamma I_t \epsilon_{t_1}^2 + \beta h_{t-1}. \tag{2.2}$$

I_t is an indicator function that equals one when ϵ_{t_1} is negative and zero otherwise.

The standard GARCH model forces a symmetric response in conditional volatility following return innovations but γ in the TARCH specification allows a asymmetric response of the conditional volatility to return innovations which means negative shocks can have a different effect on conditional volatility than positive return shocks.

Because the error term is assumed to follow a normal density, the log-likelihood function can be expressed as

$$L(\phi) = \sum_{t=1} L_t(\phi); \quad L_t(\phi) = -\log h_t - \epsilon_t^2 / 2h_t \tag{2.3}$$

For classical estimation, maximum likelihood estimation is employed. The commonly imposed inequality constraints on the parameters are $\omega > 0$, $\gamma \geq 0$, $\alpha + \gamma \geq 0$, and $\beta \geq 0$.

3. Bayesian Inference: The Random Walk Chain Metropolis - Hastings algorithm.

Bayesian inference is well-suited to impose restrictions through prior information that decribes the researcher's belief. We impose two different restrictions through the prior distribution; one is the conditional variance over time is positive and finite and the second is that conditional variance is positive, finite, and stationary. The prior distribution on the parameters is taken to be a product of independent standard normal distributions on the regression coefficients. This prior is informative, but given the prior variances' size relative to likely parameter values most of the information is due to the truncation that imposes the conditional variance restrictions. The prior distributions with positivity constraints can be written as

$$p(\mu, \lambda, \omega, \alpha, \gamma, \beta) = I(h_t) \cdot N_\mu(0,1) \cdot N_\lambda(0,1) \cdot N_\omega(0,1) \cdot N_\alpha(0,1) \cdot N_\gamma(0,1) \cdot N_\beta(0,1). \tag{3.1}$$

The second prior distribution which imposes positivity and stationarity constraints is

$$p(\mu, \lambda, \omega, \alpha, \gamma, \beta) = S(h_t) \cdot I(h_t) \cdot N_\mu(0,1) \cdot N_\lambda(0,1) \cdot N_\omega(0,1) \cdot N_\alpha(0,1) \cdot N_\gamma(0,1) \cdot N_\beta(0,1), \tag{3.2}$$

where an indicator function $(I(h_t))$ equals zero when a generated variance is negative or goes to infinity and one otherwise and another indicator funtion $(S(h_t))$ equals one when conditional variance is stationary and zero otherwise. The stationarity condition for TARCH is $\alpha + \beta + \frac{1}{2}\gamma < 1$ (Ling and McAleer, 2002). The likelihood function is the same as that for classical inference.

Since the posterior density is proportional to the likelihood function times the prior distribution and the prior has truncation due to the inequality restrictions, it is difficult to compute analytically. The Random Walk Chain Metropolis-Hastings algorithm (Koop, 2003) is used to estimate the posterior distribution. Formally, the Random Walk Chain Metropolis-Hastings algorithm generates candidate draws according to

$$\theta^* = \theta^{(s-1)} + z, \tag{3.3}$$

where z is called the increment random variable and $\theta^{(s-1)}$ is the previous parameter draw. Each draw is accepted with probability

$$\alpha(\theta^*|\theta^{(s-1)}) = \min\Big[\frac{p(\theta = \theta^*|y)}{p(\theta = \theta^{(s-1)}|y)},\ 1\Big], \tag{3.4}$$

where $p(\theta|y)$ is the posterior distribution. If a draw is rejected, the previous draw is reused.

A common and convenient choice of density for z is the multivariate normal distribution. The $var(\hat{\theta}_{ML})$ $= \widehat{\Sigma}$ from maximum likelihood estimation is used in the candidate generating density along with c, a tuning constant to adjust the acceptance rate. The candidate generating density then can be written as

$$q(\theta|\theta^{(s-1)}) = f_N(\theta|\theta^{(s-1)}, c \cdot \widehat{\Sigma}) \tag{3.5}$$

There is no general rule for an optimal acceptance rate. Koop (2003) suggests that the acceptance rate should be roughly 0.5. We set $c = 0.05$, making the acceptance rate approximately 0.45.

We generate 50,000 draws to estimate the model parameters. The posterior mean is estimated by the average of the draws. For any function of the parameters, $g(\theta)$, the posterior mean is estimated as

$$\widehat{g_S} = \frac{1}{S}\sum_{s=1}^{S} g(\theta^{(s)}). \tag{3.6}$$

4.Empirical results.

With similar data and classical inference imposing restrictions on individual parameters, the empirical finding of Glosten, Jagannatha, and Runkle (1989) is that the risk-return tradeoff is negative. Table 1 presents the results from Bayesian inference. The risk return tradeoff (represented by the parameter λ) shows a positive relation as theory predicts. This is opposite to the evidence in Glosten et al. (1989). As is common, the R^2 is very low in GARCH-M models of stock returns.

In panel A, we impose positivity constraints through the prior in (3.1) and the evidence favors a positive risk return tradeoff. However, the conditional volatility is not stationary. Panel B shows the model estimates with positivity and stationarity restrictions imposed through the prior in (3.2). The relation between market risk premium and returns in Panel B is stronger than that in Panel A.

The probabilities that the risk return tradeoff is positive are also provided. The probability of a positive tradeoff is 0.84 in panel A but it increases when stationarity is also imposed, rising to 0.98 in panel B.

The posterior distributions of the risk return tradeoff (λ) from the two models are shown in Figure 1. Figure 2 shows the news impact curve of each model. The news impact curve examines the implied relation between ϵ_{t-1} and h_t, holding constant the information dated $t-2$ and earlier. We use the form introduced by Engle and Ng (1993). The equations for the TARCH news impact curve are

$$h_t = A + \alpha \cdot \epsilon_{t-1}^2, \quad \text{for} \;\; \epsilon_{t-1} > 0 \text{ and}$$

$$h_t = A + (\alpha + \gamma) \cdot \epsilon_{t-1}^2, \;\; \text{for} \;\; \epsilon_{t-1} < 0,$$

where A $\equiv \omega + \beta \cdot \sigma^2$, σ is the unconditional return standard deviation. The figure clearly shows asymmetry in response to return innovations with the asymmetric response stronger when stationarity is imposed.

5.Conclusion.

Evidence on the sign of the risk-return tradeoff (λ) is strongly influenced by the restrictions placed on the model. It is important to restrict the model to meet basic theoretical standards (variance should be positive), but if volatility restrictions are imposed indirectly through individual parameters, it can have unexpected effects on other parameter estimates. Earlier work on U.S. stock excess returns using estimates with restricted parameters sometimes found a negative risk-return tradeoff. Here, placing restrictions directly on volatility, we find the expected positive risk-return tradeoff. Lundblad (2007) suggested small samples were causing difficulty in estimating the risk-return tradeoff. We offer an additional explanation that empirical results at odds with theory may be caused by faulty methods of imposing restrictions.

References

Baillie, R. and R. P. Degennaro, 1990. Stock Returns and Volatility. The Journal of Financial and Quantitative Analysis 25, 203-214.

Bollerslev, T., 1986. Generalized autoregressive conditional heteroskedasticity. Journal of Econometrics 31, 307-327.

Bollerslev, T., R. Engle, and J. M. Wooldridge, 1988. A Capital Asset Pricing Model with Time-varying Covariances. Journal of Politicla Economy 96, 116-131.

Campbell, J. and L. Hentschel, 1993. No news is good news: an asymmetric model of changing volatility in stock returns. Journal of Financial Economics 31, 281-381.

Engle, R., 1982. Autoregressive conditional heteroskedasticity with estimates of the variance of U.K. inflation. Econometrica 50, 987-1008.

Engle, R., D. M. Lilien, and R.P. Robins, 1987. Estimating time varying risk premia in the term structure: The ARCH-M model. Econometrica 55, 391-407., Journal of Finance 48, 1749-1778.

Engle, R. and V. Ng, 1993. Measuring and Testing the Impact of News on Volatility. Journal of Finance 48, 1749-1778.

French, K., G. Schwert,, and R. Stambaugh, 1987. Expected stock returns and volatility. Journal of Financial Economics 19, 3-30.

Geweke, J., 1989. Exact predictive densities for linear models with ARCH disturbances. Journal of Econometrics 40, 63-86.

Glosten, L., R. Jagannathan, and D. Runkle, 1989. On the relationship between the expected value and the volatility of the nominal excess return on stocks. Journal of Finance 48, 1779-1801.

Koop, G.,2003. Bayesian Econometrics.

Ling, S. and M. Mcaleer, 2002. Stationarity and the existence of moments of a family of GARCH processes. Journal of Econometrics 106, 109-117.

Lundblad, C., 2007. The risk return tradeoff in the long run: 1836-2003. Journal of Financial Economics 85, 123-150.

Nelson, D., 1991. Conditional heteroskedasticity in asset returns: A new approach. Econometrica 59, 347-370.

Rabemananjara, R. and J. M. zakoian, 1993. Threshold Arch Models and Asymmetries in Volatility. Journal of Applied Econometrics 8,31-49.

Sentana, E., 1995. Quadratic ARCH models. Review of Economic Studies 62, 639-661.

Table 1. The results of GARCH-M with TARCH specification

	Panel A		Panel B	
	Positivity		Positivity & Stationary	
Coefficient	Posterior mean	Posterior std. dev.	Posterior mean	Posterior std. dev.
$\mu \times 10^2$	0.5543	0.2203	0.2679	0.2289
λ	0.4311	0.4594	1.2452	0.4618
$\omega \times 10^3$	0.2011	0.0666	0.3903	0.0852
α	0.1536	0.0515	0.0480	0.0326
γ	0.1853	0.0872	0.2562	0.0734
β	0.8381	0.0231	0.8131	0.0317
Acceptance	0.4453		0.4402	
R^2	0.0026		-0.0035	
Prob (λ > 0)	0.8386		0.9841	

Figure 1. The posterior distribution of risk return tradeoff

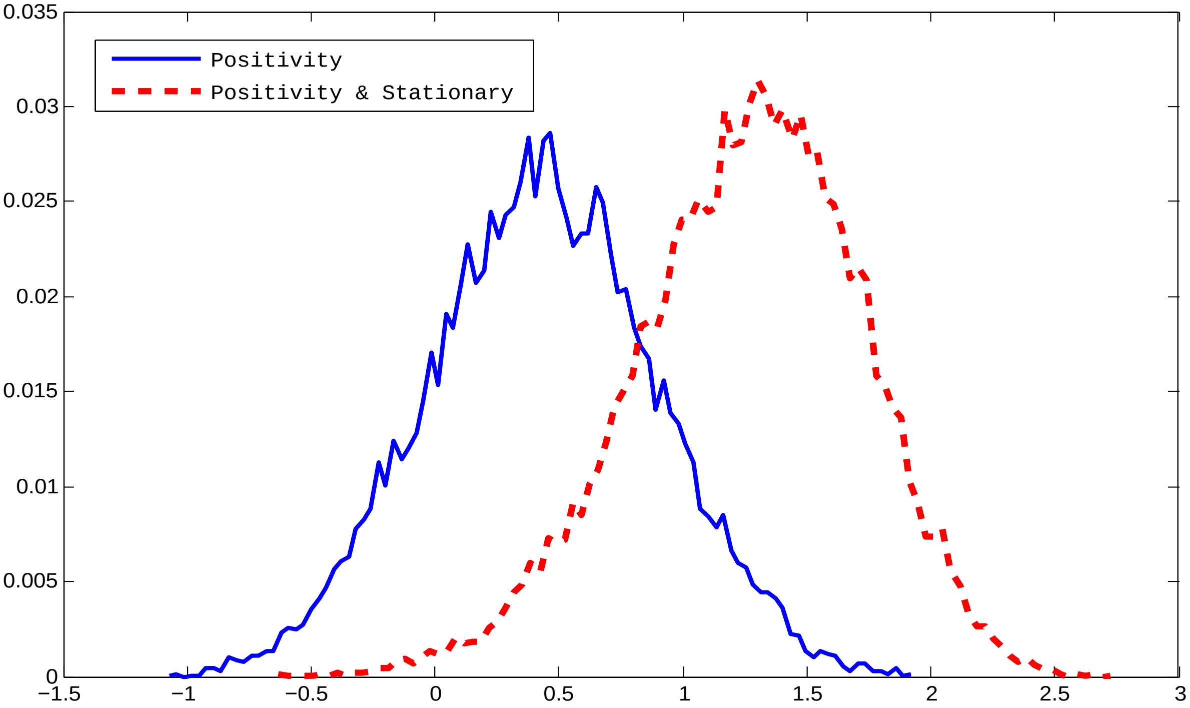

Figure 2. The News impact curves of both TARCH models