FORECASTING THE VOLATILITY OF HEALTH CARE STOCKS

Bradley T. Ewing*
Jamie B. Kruse
Mark A. Thompson

Abstract: This paper compares the performance of two popular volatility models for making out-of-sample forecasts of health care stock market volatility. We focus on portfolios designed to represent the three profit-motivated groups: providers of services, manufacturers of products, and third-party payors. Our results are consistent with Gokcan (2000) in the sense that the risk/return trade off inherent in the changing health care industry may be similar to the risk/return trade off of emerging markets. Implications for linear vs. nonlinear volatility models are discussed.

JEL Classifications: C53 Forecasting and Other Model Applications, I11 Analysis of Health Care Markets

Keywords: forecasting, volatility, GARCH models, health care markets

---

* Corresponding Author: Bradley T. Ewing, Area of ISQS, Rawls College of Business, Texas Tech University, Lubbock, Texas 79409-2101, phone: (806) 742-3939, email: bradley.ewing@ba.ttu.edu. We thank Jennifer Burgett of Standard & Poors' Corporation for her help in assembling the data set.

INTRODUCTION

Health care now accounts for a large portion of the US economy. In 2001, U.S. Health expenditures as a percent of GDP, exceeded 14%. In the same year, medical care spending represented 18.2% of total personal consumption expenditures.[1] Thus it is not surprising that institutional and individual investors should concentrate on the sectors that comprise this industry. Although the government and non-profit agencies are an important part of the U.S. health care system we can identify four other important groups: (1) patients, (2) physicians, hospitals, and medical service companies, (3) pharmaceuticals, medical supplies, and other healthcare-related products, and (4) third-party payors. Unlike patients; however, the other three groups consist of profit-maximizing entities whose goal is to increase shareholder wealth and firm value. Thus, from an investment standpoint, knowledge about how the returns of companies comprising these sectors behave over time should be particularly important. With respect to financial risk management and decision-making, a variety of techniques may be employed to forecast asset return volatility [Diebold (2001)]. In fact, the use of empirical volatility models is now a common tool for assessing risk and making reallocation decisions. The purpose of this paper is to compare the performance of two popular volatility models for making out-of-sample forecasts of health care stock market volatility. In particular, we focus on portfolios designed to represent the three profit-motivated groups. We use the term *provider* to designate a portfolio consisting of physicians, hospitals, and medical service companies, *product* denotes a portfolio of pharmaceuticals, medical supplies, and other healthcare-related products companies, and *payor* denotes a portfolio of third-party payors.

A number of different volatility models are discussed in the empirical finance literature.[2] The most popular volatility model is the generalized autoregressive conditional heteroscedasticity (GARCH) model due to its successes in modeling and forecasting return volatility. In addition, the GARCH model is capable of handling the excess kurtosis, a common feature in asset return volatility. Unfortunately, another prominent feature in asset return volatility is skewed time series, which the GARCH model cannot address. However, the exponential GARCH (EGARCH), which is a non-linear GARCH model, can handle seriously skewed time series and thus may be more appropriate in modeling and forecasting financial time series.

Franses and van Dijk (1996) discussed a number of the non-linear volatility models and their success in forecasting volatility in developed stock markets. They find that (in most cases) the non-linear GARCH models outperform the linear GARCH model. Gokcan (2000) extended Franses and van Dijk (1996) to forecast volatility in developing (or emerging) stock markets and found that the GARCH model outperformed the EGARCH model for developing stock markets.

This paper extends the Frances and van Dijk (1996) and Gokcan (2000) papers using the GARCH and EGARCH models to forecast volatility of three sector-specific stock indices within the health care industry. Since the health care industry

1 Source: Survey of Current Business, 2002, www.bea.gov/bea/dn/nipaweb/.

2 See Engle and Patton (2001) for discussion of good volatility models.

has drastically changed in structure over the past decade due to managed care [Gaynor and Haas-Wilson (1999)], examining the behavior of different health care indices using only one volatility model may be inappropriate. Furthermore, the additional information specific to the different sectors may be valuable to financial market participants and portfolio managers on selecting appropriate risk management strategies. Therefore, we compare the out-of-sample forecasting performance on three health care stock market indices using the GARCH and EGARCH volatility models. The results provide further insights into modeling and forecasting health care equities.

2. Description of Data

We use daily observations of closing prices for three Morgan Stanley Healthcare Indices from May 9, 1996 through October 28, 1999 for a total of 876 observations.[3] The three health care indices (i.e., *payor*, *provider,* and *product*) are designed to represent the managed health care organizations, health care providers, and producers of drugs and supplies sectors. The *payor* index is a portfolio of stocks that include third party payors, managed health care and health services companies such as Aetna, CIGNA, and Coventry. The nature of this sector is competitive. In fact, the competition is sometimes so fierce that it is not uncommon for a number of these firms to have experienced varying degrees of financial troubles [Dranove (2000)]. Likewise, the *provider* index represents firms in the hospital and physician services industry and is comprised of stocks from hospital management companies such as Columbia/HCA Healthcare, PhyCor, and Tenet Healthcare. The contractual arrangements between the payors and the providers due to the rise in managed care provide evidence of the inherent risk embedded in these indices [White (1999)]. The *product* index is comprised of stocks from pharmaceutical and biotechnology companies such as Pfizer, Johnson & Johnson, and Bristol-Myers Squibb. This sector is probably the most diversified of the three in that its products are also sold through standard retail distribution channels (e.g., over-the-counter drugs).

These different health care sectors may behave differently from one another due to different market characteristics; therefore, one volatility model for one sector may forecast better than another model for a different sector. However, each index consists of value maximizing firms that are therefore profit-driven enterprises. In addition, many of these firms are relatively new, at least compared to many of the firms that make up the Dow Jones Industrial Index and other manufacturing indices. Moreover, the health care industry is subject to many changes, both institutional and technologically. These factors may make the health care sectors appear more similar to emerging markets (or developing stock markets) where similar risk/return characteristics prevail. It is interesting to note that Gokcan (2000) found the linear GARCH model outperformed the nonlinear (EGARCH) model for the case of emerging stock markets while Frances and Van Dijk (1996) found just the opposite for the case of developed markets.

[3] See appendix for a detailed description of the Morgan Stanley Healthcare stock indices.

Consequently, it may be that a volatility model that forecasts well for emerging stock markets may also provide better forecasts for the health care sectors than models that have been shown to do well in otherwise developed stock markets.

Table 1 reports the summary statistics for the returns of *payor*, *provider*, and *product* health care indices. The mean daily returns range from -0.03% (*provider*) to 0.11% (*product*). However, the key features of note in Table 1 are excess kurtosis and negative skewness for all three health care return series. Since our purpose is to evaluate the out-of-sample forecasting performance of different volatility models, we exclude the last 120 days as our "hold-out" sample. In the next section, we discuss in detail the empirical methodology used to model and forecast volatility for the three health care stock market indices.

3. Empirical Methodology

The return series for each health care index is modeled as an autoregressive moving average (ARMA) process. The ARMA model may be written as:

$$(1) \qquad \phi(B)r_t = \theta(B)\varepsilon_t + \mu$$

where ϕ and θ are polynomials in the backwards shift (or lag) operator *B*, and m is a constant term. The best-fitting ARMA representation for each return series was determined by examination of the autocorrelation functions, AIC, and standard Box-Jenkins techniques.[4] The specifications were ARMA(0,2) for *payor* returns, ARMA(0,1) for *provider* returns, and ARMA(0,0) for *product* returns. Diagnostic tests indicated that each of the models was free of autocorrelation. A common feature of financial time series is that they exhibit volatility clustering. That is, periods of heightened volatility are clustered together as are periods of relative tranquility. When this feature is present, the series may have a nonconstant variance and ordinary least squares estimation would be inappropriate. Each model specification was therefore tested for the presence of autoregressive conditional heteroscedasticity (ARCH) using the Lagrange multiplier test [Engle (1982)]. It was determined that each return model exhibited ARCH effects. This finding suggests that the variance of the equation is not constant and that the variance changes over time in a way that depends on how large the errors were in the past. The Lagrange multiplier test statistics used to test for the presence of ARCH effects were significant at the 1% level or less.[5] Fortunately, Engle (1982) developed a class of models capable of handling this type of situation.
In order to examine healthcare returns with time-varying volatility, we estimate the mean and variance of the return series simultaneously via the method of maximum

[4] See the time series econometrics texts of Harvey (1994) for an overview of the Box-Jenkins method for specifying an ARMA model.

[5] Results of the ARCH LM test are available upon request.

[6] The (1,1) specification was chosen in a manner similar to that which was used to determine the order of the ARMA models. Moreover, Bollerslev, et al. (1992) conclude that the GARCH(1,1) model is preferred in most cases.

likelihood. To measure the volatility, we consider the GARCH(1,1) and EGARCH(1,1) models.[6] The (conditional) variance equation for the GARCH(1,1) model is given by:

$$(2) \qquad h_t^2 = c + \alpha\varepsilon_{t-1}^2 + \beta h_{t-1}^2$$

where $V\ (\varepsilon_t|\Omega_{t-1}) = h_t^2$ is the conditional variance of ε_τ with respect to the information set $\Omega_{\tau-1}$. We computed the quasi-maximum likelihood covariances and standard errors as described in Bollerslev and Wooldridge (1992). The model is estimated under the assumption that the errors are conditionally normally distributed. Both α and β are non-negative constant parameters and $\alpha+\beta<1$. These restrictions on the parameters prevent negative variances (Bollerslev, 1986).

A feature of the GARCH model is that forecasts of future variance are linear in current and past variances [Campbell, Lo, and MacKinlay (1997)]. This implies that the GARCH model will treat positive and negative shocks symmetrically. That is, the sign of a shock is independent of the effect that it has on conditional volatility. However, Engle and Ng (1993) found that negative shocks to returns had a much greater effect on the volatility of returns than did positive shocks. In this sense, it is often said that "bad" news (i.e., an unexpected fall in stock prices) has a more pronounced effect on volatility than does "good" news. Nelson's (1991) *exponential* GARCH (or EGARCH) model is capable of detecting whether positive and negative shocks have an asymmetric impact on volatility. The (conditional) variance equation for the EGARCH(1,1) model is given by:

$$(3) \qquad \log\left(h_t^2\right) = \omega + \alpha\left|\varepsilon_{t-1}/h_{t-1}\right| + \delta\left(\varepsilon_{t-1}/h_{t-1}\right) + \beta\log\left(h_{t-1}^2\right)$$

Note that the left-hand side of (3) is the log of the conditional variance and guarantees that forecasts of conditional variance are non-negative. While Nelson (1991) assumes a generalized error distribution, we simply assume normally distributed errors. The estimates will be the same under either assumption except for the intercept term in our model, which will differ from that of the original Nelson model by $\alpha(2/\pi)^{-.5}$. In the EGARCH model, the coefficient on $|\varepsilon_{t-1}/h_{t-1}|$, α, represents what is referred to as the "size" effect. The size effect provided information as to the degree to which volatility is affected by a shock regardless of whether it was negative or positive. The "sign" effect is given by δ. *If* $\delta \neq 0$, then there are said to asymmetric effects to volatility. In particular, when δ is negative and statistically significant, one may conclude that a negative shock increases volatility more than a positive shock of equal magnitude. The β coefficient reveals information as to the degree of volatility persistence. The larger this coefficient is, the greater the degree of persistence in so much as a shock's affect on volatility will take more time to dissipate.

For the purpose of assessing the performance of these GARCH-class models in forecasting the volatility, we construct a measure of true volatility to compare with the forecasted volatility from the GARCH and EGARCH models [see Pagan and Schwert (1990), Day and Lewis (1992), Franses and van Dijk (1996); Gokcan (2000)]. Similarly, we define the true unconditional volatility as:

$$(4) \qquad \sigma_t^2 = \left(r_t - \bar{r}\right)^2$$

where s^2 is the unconditional volatility, r_t is the actual daily return for day t, and $\bar{r}$ is the expected return for day t. The squared difference between actual returns and the (moving) average returns gives the implied volatility as noted in equation (4). The expected return for 5/11/99 is measured by calculating the arithmetic average of daily returns from 5/9/96 to 5/10/99; the expected return for 5/12/99 is measured by calculating the arithmetic average of daily returns from 5/10/96 to 5/11/99. This is repeated for the next 118 days. That is, we estimate the GARCH and EGARCH models using 755 observations and saving the last 120 observations for out-of-sample forecasting comparison. To obtain the one-period-ahead forecasting errors for the two volatility models, we used the following equation:

$$(5) \qquad u_{t+1} = \sigma^2_{t+1} - \hat{h}^2_{t+1}$$

where σ^2 is the "true" unconditional volatility h^2_{t+1} and is the forecasted variance generated from the GARCH(1,1) and EGARCH(1,1) models.

4. Discussion

The in-sample estimation of the GARCH and EGARCH volatility models is reported in Table 2. The parameter estimates for the GARCH and EGARCH models are all significant for the *payor* (except for the constant terms), *provider*, and *product* indices. In the EGARCH model, the sign effect given by d is negative and significant for all three health care sectors. These results are consistent with Engle and Ng (1993) and indicate that positive and negative shocks have asymmetric effects on the volatility of health care returns. The Akaike Information Criterion (AIC) was used to compare the in-sample estimation of the two volatility models. By this measure, the EGARCH(1,1) model outperforms the GARCH(1,1) model in all three health care sectors.

Given the findings of Frances and van Dijk (1996) and Gokcan (2000), it is not necessarily the case that the model that performs best in-sample will also be the model that provided the best out-of-sample forecasts of volatility. In order to assess the out-of-sample performance of the models we computed the mean squared error (MSE) associated with several different forecasting horizons. The MSE is a useful and intuitively appealing measure for choosing the best forecasting model. For a sample of size T, we have:

$$(6) \qquad MSE = T^{-1} \sum_{t=1}^{T} u_t^2$$

Table 3 reports the MSE for the 30, 60, 90, and 120 one-step-ahead forecasts generated from the GARCH(1,1) and EGARCH(1,1) models. The results indicate that the GARCH(1,1) model outperforms the EGARCH(1,1) model for the *provider* and *product* sectors. The *payor* sector produced mixed results. The results for the *payor* returns are consistent with the other health care indices in that the GARCH(1,1) model has smaller forecast error for the 30 and 90 day time horizons. However, the EGARCH(1,1) produces smaller forecasting errors than the GARCH(1,1) model for the 60 and 120 day forecast horizons.

5. Conclusions

In this paper, we have focused on forecasting return volatility for three distinct sectors of the health care industry. In particular, we compare the GARCH and EGARCH in-sample estimation as well as the out-of-sample forecasts for the health care payor, provider, and product sectors. We find that the EGARCH(1,1) model is superior to the GARCH(1,1) model in-sample, regardless of which health care sector is analyzed. However, the out-of-sample forecast evaluations indicate that the GARCH(1,1) model produces smaller forecast errors than the EGARCH(1,1) even with significantly skewed time series. The only exception is with the payor index at the 60 and 120 day horizons. These results indicate some similarity between the volatility of returns in health care and the volatility of returns in emerging markets. Is there a case to be made for viewing the sectors of the health care industry as emerging markets? The calculated inflation rates for hospital services, physician services, and prescription pharmaceuticals have all exceeded the general rate of inflation. In the 1990's cost pressures have moved more patients toward managed care. By 2001, Fee-for Service health care enrollment for covered workers fell to 7% from 73% just 13 years earlier (Folland, et al, 2004). It is clear that the health care industry is evolving with no clear indication of what the stable configuration of for-profit entities will be. Our results are consistent with Gokcan (2000) in the sense that the risk/return trade off inherent in the changing health care industry may be similar to the risk/return trade off of emerging markets. If this is the case, the linear GARCH model will likely produce superior forecasts.

REFERENCES

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

Bollerslev, T., R. Chou, and K. Kroner. 1992. ARCH modeling in finance: a review of the theory and empirical evidence. Journal of Econometrics, 52, 5-59.

Bollerslev, T. and J. Wooldridge. 1992. Quasi-maximum likelihood estimation and inference in dynamic models with time-varying covariances. Econometric Reviews, 11, 143-172.

Campbell, J., A. Lo, and A. MacKinlay. 1997. The Econometrics of Financial Markets. Princeton, NJ: Princeton University Press.

Day, T. and C. Lewis. 1992. Stock market volatility and the information content of stock index options. Journal of Econometrics, 52, 267-287.

Diebold, F. 2001. Elements of Forecasting, second edition. Cincinnati: South-Western.

Dranove, D. 2000. The Economic Evolution of American Health Care. Princeton, NJ: Princeton University Press.

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

Engle, R. and A. Patton. 2001. What good is a volatility model? Quantitative Finance, 1, 237-245.

Engle, R. and K. Ng. 1993. Measuring and testing the impact of news on volatility. Journal of Finance, XLVIII, 1749-1778.

Folland, S. A. Goodman and M. Stano. 2004. The Economics of Health and Health Care 4th ed., Upper Saddle River, New Jersey; Pearson Prentice Hall.

Franses, P. and R. Van Dijk. 1996. Forecasting stock market volatility using (non-linear) GARCH models. Journal of Forecasting, 15, 229-235.

Gaynor, M. and D. Haas-Wilson. 1999. Change, consolidation, and competition in health care markets. Journal of Economic Perspectives, 13, 141-164.

Gokcan, S. 2000. Forecasting volatility of emerging stock markets: Linear versus non-linear GARCH models. Journal of Forecasting, 19, 499-504.

Harvey, A. 1994. Time Series Models. Cambridge, MA: MIT Press.

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

Pagan, A. and W. Schwert. 1990. Alternative models for conditional stock volatility. Journal of Econometrics, 45, 267-290.

White, W. 1999. Physicians and the changing health care marketplace. Quarterly Review of Economics and Finance, 39, 439-444.

Table 1. Descriptive Statistics

	PAYOR	PROVIDER	PRODUCT
Mean	0.0200	-0.0273	0.1142
Median	0.0300	0.0756	0.1188
Maximum	8.9273	5.4399	5.3217
Minimum	-15.0404	-9.8232	-8.4412
Standard Deviation	1.6531	1.5358	1.2891
Skewness	-1.1063	-0.9238	-0.6979
Kurtosis	15.1899	8.5732	8.6827

Note: PAYOR, PROVIDER, and PRODUCT are daily returns calculated as $r_t=[log\ (P_t) - log\ (P_{t-1})]*100$ where P is the stock price of the healthcare index at time t.

Table 2. GARCH and EGARCH Parameter Estimates

	PAYOR	PROVIDER	PRODUCT
GARCH (1,1)			
Constant (c)	0.1506	0.1058[a]	0.2033[a]
Arch (a)	0.0424[b]	0.1552[a]	0.1701[c]
Garch (b)	0.9046[a]	0.8151[a]	0.7136[a]
lnL	-1430.9	-1309.3	-1215.5
AIC	3.8063	3.4817	3.2305
EGARCH (1,1)			
Constant (w)	-0.0197	-0.1389[b]	-0.0432[b]
Size effect (a)	0.0667[a]	0.2295[b]	0.0668[a]
Sign effect (d)	-0.1327[a]	-0.1038[b]	-0.1245[a]
Egarch (b)	0.9675[a]	0.9506[a]	0.9869[a]
lnL	-1397.3	-1298.3	-1200.3
AIC	3.7200	3.4551	3.1928

Notes: The superscripts a, b, c denote significance at the 1, 5, 10% level respectively. Daily returns are calculated as $r_t=[log\ (P_t) - log\ (P_{t-1})]*100$ where P is the stock price of the healthcare index at time t. The in-sample period for estimation of returns and conditional variance is from May 9, 1996 to May 10, 1999.

Table 3. Mean Squared Error Terms

Index	Model	Out-of-Sample Forecast Period			
		30 days	60 days	90 days	120 days
PAYOR	GARCH EGARCH	9.78 10.30	17.02 16.80	12.95 13.68	137.75 132.94
PROVIDER	GARCH EGARCH	5.66 6.13	7.96 8.36	6.60 7.24	8.40 8.80
PRODUCT	GARCH EGARCH	4.87 8.03	3.56 5.57	4.08 5.69	5.82 7.46

Note: Daily returns are calculated as $r_t = [log\ (P_t) - log\ (P_{t-1})]*100$ where P is the stock price of the healthcare index at time t.

Appendix

Healthcare Payor Index

Description: The Morgan Stanley Healthcare Payor Index represents a portfolio of widely held and highly capitalized managed healthcare and health services companies and is designed to provide a broad measure of this field.

Components:
Aetna Inc. (AET)
Mid Atlantic Medical Services, Inc. (MME)
CIGNA Corporation (CI)
Oxford Health Plans, Inc. (OXHP)
Coventry Corporation (CVTY)
PacifiCare Health Systems, Inc. (PHSYB)
First Health Group Corporation (FHCC)
Trigon Healthcare, Inc. (TGH)
Foundation Health Systems, Inc. (FHS)
United Healthcare Corporation (UNH)
Humana, Inc. (HUM)
Wellpoint Health Networks, Inc. (WLP)

Healthcare Provider Index

Description: The Morgan Stanley Provider Index is designed to provide a broad measure of the performance of hospital management and medical/nursing services companies and includes shares of widely held and highly capitalized companies in this field.

Components:
Apria Healthcare Group, Inc. (AHG)
Lincare Holdings, Inc. (LNCR)
Beverly Enterprises, Inc. (BEV)
Manor Care, Inc. (MNR)
Columbia/HCA Healthcare Corporation (COL)
MedPartners, Inc. (MDM)
Genesis Health Ventures, Inc. (GHV)
PhyCor, Inc. (PHYC)
Health Care and Retirement Corporation (HCR)
Quorum Health Group, Inc. (QHGI)
Health Management Associates, Inc. (HMA)
Tenet Healthcare Corporation (THC)
HEALTHSOUTH Corporation (HRC)
Vencor, Inc. (VC)
Integrated Health Services, Inc. (IHS)

Healthcare Product Index

Description: The Morgan Stanley Healthcare Product Index represents a portfolio of widely held and highly capitalized pharmaceutical, biotechnology, medical distributor and medical electronics companies, and is designed to provide a broad measure of companies in this field.

Components:
Abbott Laboratories (ABT)
Cardinal Health Inc. (CAH)
Mylan Laboratories, Inc. (MYL)
American Home Products Corporation (AHP)
Chiron Corporation (CHIR)
Pfizer, Inc. (PFE)
Amgen, Inc. (AMGN)
Genzyme Corporation (GENZ)
Pharmicia & Upjohn, Inc. (PNU)
Bard (C.R.), Inc. (BCR)
Glaxo Wellcome pic (GLX)
Schering-Plough Corporation (SGP)
Baxter International Inc. (BAX)
Guidant Corporation (GDT)
SmithKline Beecham Pic (SBH)
Becton, Dickinson & Co. (BDX)
Johnson & Johnson (JNJ)
St. Jude Medical, Inc. (STJ)
Biogen, Inc (BGEN)
Lilly (Eli) & Co. (LLY)
Stryker Corporation (STRY)
Boston Scientific Corporation (BSX)
Medtronic, Inc. (MDT)
Warner-Lambert Co. (WLA)
Bristol-Myers Squibb Company (BMY)
Merek & Company, Inc. (MRK)