Journal of Financial and Strategic Decisions
Volume 13 Number 2 Summer 2000

GARCH MODELS AND THE STOCHASTIC PROCESS UNDERLYING EXCHANGE RATE PRICE CHANGES

Ken Johnston* and Elton Scott**

Abstract

This study investigates the extent of the contribution of the original GARCH model to our understanding of the stochastic process underlying exchange rate price changes, and examines if the movement of current research to GARCH type models exclusively is warranted. GARCH(1,1) parameters are calculated on a yearly basis and used to standardize the exchange rate price change data. Frequency distributions and statistical tests indicate that independence still exists after standardization. This indicates that GARCH type models alone are inadequate since all are similar in form, and would have difficulty in accounting for such independence. It could be argued that the poor performance of the GARCH model is due to the models incorrect assumption of a normal distribution. This argument is tested by comparing the GARCH standardized data with mean -variance standardized data, which makes no assumptions about the distributional form. Results of likelihood ratio tests, question the significance of conditional volatility, in explaining exchange rate price changes. Curiously there are cases where GARCH e2(t-1) parameters are significant when tests for first-order heteroskedasticity are not significant; this suggests that the model may be misspecified. Overall, results indicate that although previous research finds that volatility clustering plays a role in determining the stochastic process, it is not the dominate factor. This study questions the contribution of the GARCH type models. We discuss implications of our results.

INTRODUCTION

Recent empirical research relating to the probability distributions of daily spot exchange rate changes use dependent models, where variance is conditional. Studies have focused on ARCH-type, GARCH processes. These models propose that volatility varies over time and is persistent. Although these papers contribute to our understanding of the nature of volatility it has yet to be shown that these models improve our ability to describe exchange rate price changes. If volatility clustering is the primary factor in explaining daily exchange rate price changes, these models would be superior to any independent models and significant independence would not exist in the data after removing the ARCH/GARCH type effects. The ARCH/GARCH models are referred to as dependent models since they do not have independent random variables. An example of an independent model is a mixture of normals, since it assumes that the observations are independent random variables.

The GARCH model is a more general case than the ARCH model. In their original form, a normal distribution is assumed, with a conditional variance that changes over time. For the ARCH model, the conditional variance changes over time as a function of past squared deviations from the mean. The GARCH processes variance changes over time as a function of past squared deviations from the mean and past variances. Formally, the GARCH model can be expressed as follows:

Equation 1

$$Y(t) = x(t)P + e(t)$$

Equation 2

$$e(t) \quad \Phi_{t-1} \sim N(0, \sigma^2(t))$$

*Georgia Southern University
**Florida State University

Equation 3

$$\sigma^2(t) = \alpha_0 + \sum_{I=1}^{p} \alpha_i e^2(t-I) + \sum_{j=1}^{q} \beta_i \sigma^2(t-j)$$

Where the conditional information set at time t-1 is denoted Φ_{t-1}. The variance of the ARCH process exists when $\Sigma\alpha_i < 1$ and is given by, $Var(Y(t)) = \alpha_0 / (1 - \Sigma\alpha_i)$. In this study $Y(t)$ is equal to the change in $log(S(t))$, the log of the spot exchange rate. $X(t)$ is a 1xk vector of lagged endogenous variables included in the information set. P is a kx1 vector of unknown parameters. p and q are the order of the process. In this study a GARCH(1,1) model is used. Equation (3) reduces to:

Equation 4

$$\sigma^2(t) = \alpha_0 + \alpha_i e^2(t-1) + ß_i\sigma^2(t-1)$$

Although the proper procedure to fit a GARCH model entails diagnostic checking in addition to estimation, We focus on the GARCH(1,1) due to the support found for this model in previous research.

Recent studies have supported an ARCH/GARCH type process. Hsieh (1988) found highly significant autocorrelations for the squared residuals of daily data, but an ARCH(12) process accounts for most of the nonlinear stochastic dependencies found. Studies by Taylor (1986), McCurdy and Morgan (1988), and Kugler and Lenz (1990), find that the GARCH(1,1) model has a superior fit than the ARCH(12) model.

Fujihara and Park (1990) used weekly futures price changes and spot price changes, to test the ARCH process, the compound normal process, the mixed jump diffusion process, and the Student t-distribution process. These generating processes assume independence, with the exception of the ARCH process. The results support the ARCH model for three out of the five currencies examined.

Bollerslev and Ghysels (1996) propose using a periodic (P-ARCH(1,1), P-GARCH(1,1)) structure to formulate time series models for conditional heteroskedasticity to capture the repetitive seasonal time variation of the second order moments. Using Monte Carlo simulation, and an empirical example with the mark/pound exchange rate, they show a loss in efficiency when this periodic factor is ignored.

Baillie, Bollerslev, and Mikkelson (1996) introduce the fractionally integrated generalized autoregressive conditional heteroskedastic (FIGARCH) processes. The influence of lagged squared innovations has a slow hyperbolic rate of decay which tends to zero for long lags. It has the resemblance to the fractionally integrated class of processes for the conditional means. It has similar flexibility in modeling the persistence of shocks to the conditional variance process. Using a daily time series of the German mark/ U.S. dollar exchange rate, they show that the FIGARCH(1,d,1) model is superior to both the GARCH(1,1) and the integrated GARCH(1,1) (IGARCH).[1] There are other forms of ARCH/GARCH processes, for a through review of ARCH/GARCH modeling in finance see Bollerslev, Chou and Kroner (1992).

A problem with the majority of these studies is that independent models are not considered. We can not determine from these studies if the focus on modifications to the original ARCH/GARCH process are a positive step. Therefore the purpose of this study is to determine if the movement of current research to GARCH type models exclusively is warranted. We also examine if conditional volatility in the form of the original GARCH model is the primary factor in the attempt to determine the stochastic process underlying exchange rate price changes. These questions will be examined using two approaches. First, the GARCH effects will be removed by standardizing the data, using the GARCH model derived from maximum likelihood estimation. Tests of independence will then be conducted on this standardized data. If a GARCH(1,1) process is the appropriate process underlying exchange rate price changes, removal of these effects will result in a single normal distribution, suggesting that independent models do not need to be considered in future research.

In our second approach, the GARCH standardized data will be compared to mean-variance standardized data to determine if GARCH standardization is an improvement over the most naive alternative. Again, if the GARCH(1,1) model is the appropriate process, it's standardized estimates should result in a series of observations that fit a one normal distribution better than the data that has been standardized by using just the mean and variance of the observations.

There are many specifications for the GARCH model, all are not examined here. But if independence still exists after standardization, and mean-variance standardization is found to be superior, this puts most GARCH type models into question since all are similar in form, being dependent models with conditional variance.

METHODOLOGY

The objective of determining the stochastic process underlying daily exchange rate price changes is to allow development of better exchange rate pricing models. Empirical evidence indicates that parameters for the models shift over time (See Johnston, and Scott (1999)), therefore it is more appropriate to calculate model parameters on a yearly basis than over long time horizons as in previous research. Accurate descriptions of the short term distributions would allow for development of improved exchange rate pricing models. In this study, the parameters of the GARCH(1,1) model are calculated on a yearly basis, using maximum likelihood estimation (IMSL subroutine DNCONG). The estimated parameters for that year and the previous year are used to standardize the data. This way, both in sample and out of sample are tested.

By standardizing the data with the mean and the conditional variance of the GARCH model, the GARCH effects are eliminated from the data.

Maximum likelihood estimation is then conducted on this standardized data examining one normal, and a mixture of two normals. Two normals are examined, to determine if independence still exists in the data. Once the maximum likelihood parameter estimates are obtained, the standardized data is assigned to one of the two normals based on a maximum likelihood approach. The classification rule is to select the distribution j for each observation t that has the largest posterior probability, that is:

Equation 5

$$Max_j \Gamma_j p(S(t)) \quad (\alpha_j, \sigma^2_j)$$

Where Γ_j is the sample proportion of the total observations from distribution j, and $p(S(t)/(\alpha_j,\sigma^2_j))$ is a normal density function with α_j mean and variance σ^2_j.

Frequency distributions are calculated for the raw data, and the standardized GARCH(1,1) data. The data is also standardized by using the annual, and previous years mean and variance. The maximum likelihood approach discussed above is applied. They are calculated to determine if the GARCH model frequency distributions are an improvement over the raw data's and the most naive standardization alternative.

Standardization will bring in the tails of the overall data distribution. If it significantly removes the impact of the second normal distribution, which could be described as the jump distribution (due to abnormal information), then independent models are not the direction to go in future research. If independence is still indicated then the movement of the research to purely dependent models is inappropriate.

With the standardized data assigned to the two normal distributions, the hypothesis that the two groups of observations have identical distributions can be tested, using the Kolomogorov-Smirnov statistic, hence test for independence. If we find that the standardized GARCH data has a mixture of two normal distributions, although it says nothing about the validity of the conditional variance component of the GARCH model, it does question the validity of the GARCH class of models since these models have difficulty in accounting for such independence.

Likelihood ratio tests are conducted on the maximum likelihood estimates for the normal distributions from the mean-variance, and GARCH standardized data. The exact data sample in each case is not the same since we are standardizing with different parameters. The likelihood ratio test is still appropriate, the raw data is the same in every test and the standardization can be thought of as different restrictions put on the model. The likelihood ratio test is $2(L_{mv} - L_g)$, where L_g is the GARCH standardized maximum log likelihood estimates and L_{mv} is the mean-variance standardized maximum log likelihood estimates. This test has a X^2 distribution with k degrees of freedom where k is the number of restrictions being tested. One would naturally speculate that if the GARCH coefficients are significant, then the likelihood value of the GARCH standardized data would be larger than that of a constant mean and variance standardized data. Findings to the contrary, again question the validity of the GARCH model.

DATA

The exchange rate data consists of the daily closing spot prices for the British pound, Canadian dollar, German mark, and Japanese yen versus the U.S. dollar from the Chicago Mercantile Exchange for the years 1978 to 1987, and from the Merrill Lynch debt markets group's fixed income research data base for the years 1988 to 1992. The daily series represents changes between business days with no adjustment for holidays.

RESULTS

Table 1 shows the yearly significance of the GARCH(1,1) parameters. Both GARCH(1,1) conditional variance parameters are significant in only 17 of the possible 60 cases. The Canadian dollar is the only currency where there is 5 years of consecutive data with both GARCH parameters significant (84,85,86,87,88). The study will focus on this Canadian data since it gives the GARCH model the highest probability of success.[2]

TABLE 1
Yearly Significance of the GARCH(1,1) Parameters
(1987-1992)

	None	e^2_{t-1} only	σ^2_{t-1} only	Both Parameters	Years Both Parameters Significant
British	5	4	3	3	79,80,92
Canada	4	2	1	8	78,82,84,85,86,87,88,90
German	7	3	1	4	80,83,85,87
Japan	7	2	4	2	79,84
Total	23	11	9	17	

Table 2 reports the results of first order heteroskedasticity tests for each year where the GARCH parameters are significant. Lagrange multiplier (LM) test and the Portmanteua Q-test, test statistics are used to test for first order heteroskedasticity. Out of the 17 cases where both GARCH(1,1) conditional variance parameters are significant, there are 7 cases where the tests for first order heteroskedasticity are insignificant. There is no first order heteroskedasticity found in British 92, Canada 82, 87, 88, 90, and German 80, 87 data.

Casual empiricism might suggest that the elimination of heteroskedasticity may be due to GARCH-type models being used in practice, making the market more efficient and removing the heteroskedasticity. The problem with this line of reasoning is that if the use of the GARCH models removed the first-order heteroskedasticity, the $e^2(t\text{-}1)$ parameters of the GARCH(1,1) models should no longer be significant in these periods. What is this GARCH model first-order parameter explaining if there is no first-order heteroskedasticity? They may be picking up higher order heteroskedasticity, if that is the case the models used here are misspecified. It could be argued that GARCH effects exist in the data, but outliers are causing the tests for first-order heteroskedasticity to be insignificant. The question then becomes why do outliers play a more significant role in the tests for first order heteroskedasticity than they do in the GARCH modeling process. It could be caused by differing assumptions between the two tests and the model. The original GARCH model assumes a normal distribution with conditional variance that changes over time. The LM and Q statistics are computed from OLS residuals assuming the disturbances are white noise. The Q and LM statistics have an approximate X^2 distribution under the white noise hypothesis. The X^2 distribution is skewed positively but as the degrees of freedom increases it approaches the shape of a normal distribution. Since our degrees of freedom in both tests are over 200, we believe that it is not the assumptions that are causing the differences found between the tests for first-order heteroskedasticity and the significance of the parameter for $e^2(t\text{-}1)$ in the GARCH(1,1) model.

Until these questions are answered and we are confident the GARCH models are explaining what they are theoretically designed to, the movement to more complex GARCH type models is premature.

Table 3 reports the sample statistics for the Canadian overall and standardized data. The mean and standard deviations increase with standardization. The standardized data now measures how many standard deviations a data point is from the mean, the data has been scaled. The skewness and kurtosis is larger in each case for the 2nd normal distribution, which is expected, the 2nd normal is picking up the outliers. The kurtosis of the data reduces significantly when the data is separated into 2 distributions. The skewness increases with standardization for all cases (overall, normal 1, and normal 2). Again, this is due to the standardization or scaling of the data. Also shown are the results of the normality tests for the various distributions. The Shapiro - Wilk W statistic is used.[3] These

TABLE 2
First Order Heteroskedasticity Tests
In Years Where both Conditional Variance
Parameters are Significant
(1987-1992)

	Q	Prob>Q	LM	Prob>LM
British				
1979	7.8986*	.0049	7.8991*	.0049
1980	17.2327*	.0001	17.0225*	.0001
1992	.0059	.9388	.0079	.9291
Canada				
1978	19.1527*	.0001	18.9096*	.0001
1982	.1539	.6948	.1599	.6892
1984	43.1682*	.0001	42.7215*	.0001
1985	58.8913*	.0001	58.2389*	.0001
1986	17.5788*	.0001	17.4217*	.0001
1987	1.8059	.1790	1.7875	.1812
1988	1.9495	.1626	1.9405	.1636
1990	.8144	.3668	.8048	.3697
German				
1980	2.3866	.1224	2.3757	.1232
1983	8.7477*	.0031	8.5565*	.0034
1985	2.9727*	.0847	2.9706*	.0848
1987	2.2630	.1325	2.2882	.1304
Japan				
1979	4.2382*	.0395	4.1350*	.0420
1984	25.2111	.0001	25.0247	.0001

*Significant for alpha < .1
Q= Portmanteau Q-test
LM = Lagrange multiplier test

results indicate that the distributions are not normally distributed. The hypothesis that the data comes from a normal distribution is rejected in all cases at the .01 level of significance. Therefore, the non-normality found in independent models is not caused by a failure to take into account the GARCH like dependence. When the data is standardized by the GARCH model, the resulting distributions fail the normality test. It could be argued that any evidence found in this study against the GARCH models, results from the GARCH model's incorrect assumption of a normal distribution. This argument can be tested by comparing the GARCH standardized data to the mean-variance standardized data. Even if the actual distribution is nonnormal as long as conditional volatility plays a significant role, the GARCH standardized data should fit a one normal distribution better than the mean-variance standardized data. The mean-variance standardized data is not making any assumptions about the distributional form. By definition the GARCH(1,1) standardized data comes from one normal. Therefore when we fit the GARCH standardized data to one normal distribution, it's fit should be superior. Findings to the contrary put the significance of conditional variance into question regardless of the actual distributional form.

Figures 1, 2,3, show the in sample frequency distributions for the Canadian raw data, Canadian mean-variance standardized data, and the Canadian GARCH standardized respectively.[4] Frequency distributions before and after standardization indicate that independence exists in the data that can not be accounted for by the GARCH model. The second distribution still contains a significant number of data points in every case. Figure 3, the GARCH standardized second distribution contains 208 (17%) of all the data.

Frequency distributions, are examined for all currencies (Overall, 2 normals) in 5 year intervals (1978-1982, 1983-1987, 1988-1992). The results are similar to the Canadian 5 year interval in this study, two distinct distributions. This indicates that the conclusions drawn are not time period or currency specific.

TABLE 3
Summary Statistics and Normality Tests for the Canadian Dollar. In Sample (1984-1988) Out of Sample (1985-1989)

	Mean	σ	Skewness	Kurtosis	W:Normal	Prob<W
Unstandardized						
Overall (n = 1254)	.000036	.003137	.01367	9.67571	.91728	.0001
Normal 1 (n = 1137)	.000083	.001887	.01759	-.01768	.97806	.0001
Normal 2 (n = 117)	-.000420	.008437	.16362	-.38936	.92897	.0001
Mean-Var Standardized						
A) In SampleOverall (n = 1254)	-.001070	1.00263	-.02503	8.16449	.91527	.0001
Normal 1 (n = 1216)	.021001	.75500	.01999	.67674	.97527	.0001
Normal 2 (n = 38)	-.707380	3.84656	.54413	-1.21197	.85514	.0001
B) Out of Sample Overall (n = 1262)	.018377	1.06299	.06394	16.52215	.90365	.0001
Normal 1 (n = 987)	.050389	.46405	-.05096	-.82319	.96057	.0001
Normal 2 (n = 275)	-.096520	2.09962	.19782	2.98018	.94523	.0001
GARCH Standardized						
A) In SampleOverall (n = 1254)	-.04781	.93335	-.45789	4.20539	.95616	.0001
Normal 1 (n = 1046)	-.15480	.55912	-.47021	-.04963	.94855	.0001
Normal 2 (n = 208)	.49023	1.82928	-1.13502	.66715	.81764	.0001
B) Out of SampleOverall (n = 1262)	-.01503	.96195	-.46103	6.86617	.95187	.0001
Normal 1 (n = 1235)	.01440	.78948	.10449	.36920	.98163	.0001
Normal 2 (n = 27)	-1.36135	3.65710	.69843	-.79363	.84439	.0001

*Significant at the .01 level

Table 4 reports the results of the independence tests. To test for independence we use the non-parametric Kolomogorov-Smirnov statistic. This allows to test for independence without assuming any specific distributional form. The null hypothesis is that the two groups of observations have identical distributions. The test statistic is:

Equation 6

$$D = Max_t|F_1(x_t) - F_2(x_t)|$$

Where $F_1(x_t)$ is the observed cumulative relative frequency for the t^{th} data value from distribution *1* and $F_2(x_t)$ is the observed cumulative frequency for the t^{th} data value from distribution *2*. If D is greater than its critical value, the null hypothesis is rejected, thus supporting independence.

FIGURE 1
Frequency Distributions
Canadian Raw Data

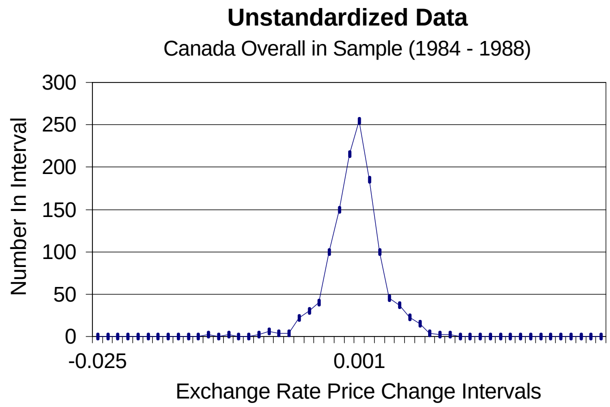

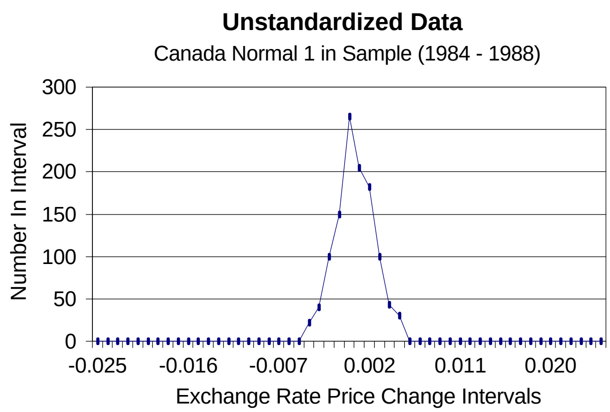

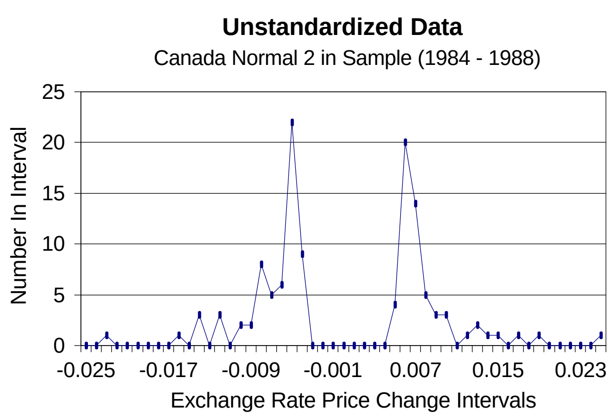

Sample Sizes: Overall = 1254 Normal 1 = 1137 Normal 2 = 117

FIGURE 2
Frequency Distributions
Canadian Mean-Variance Standardized

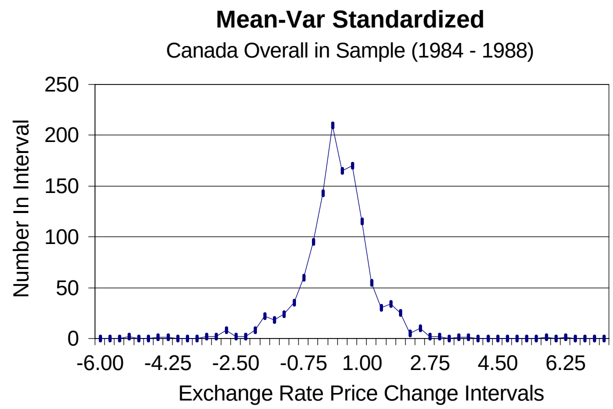

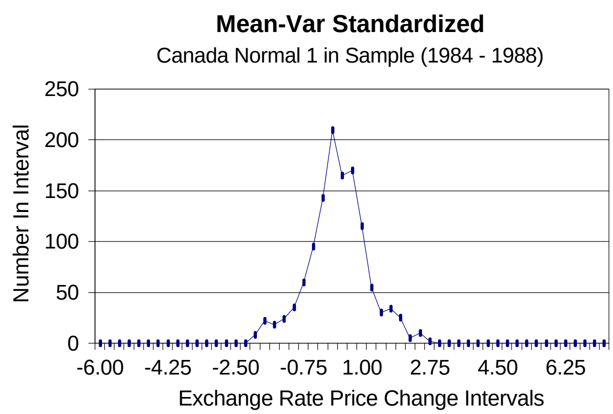

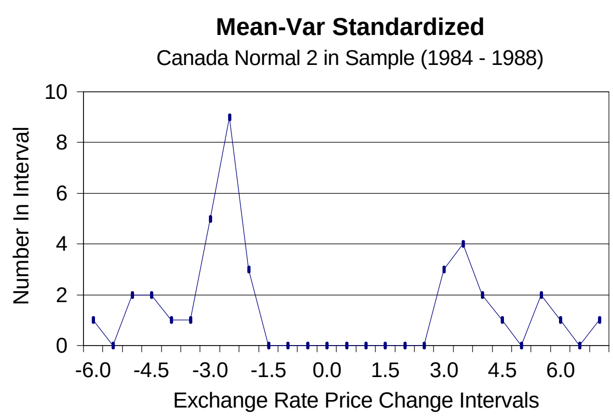

Sample Sizes: Overall = 1254 Normal 1 = 1216 Normal 2 = 38

FIGURE 3
Frequency Distributions
Canadian GARCH Standardized

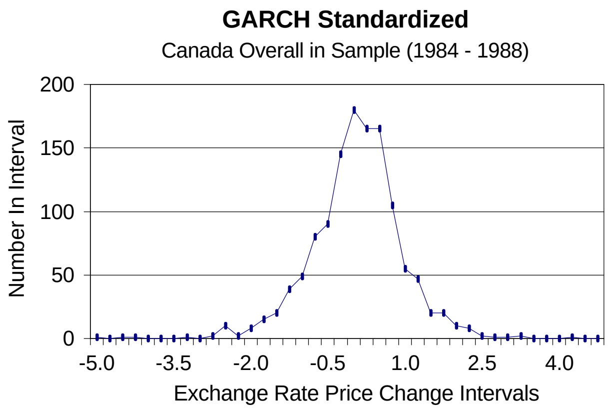

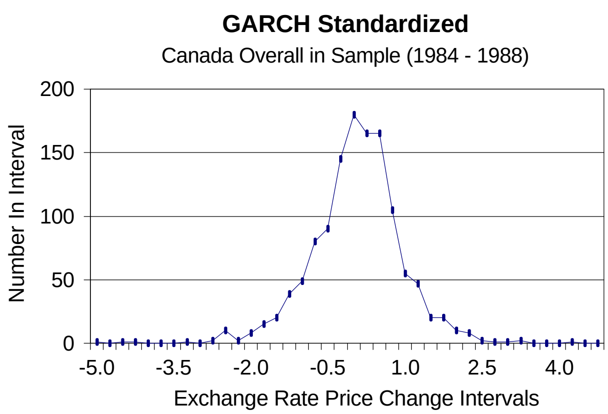

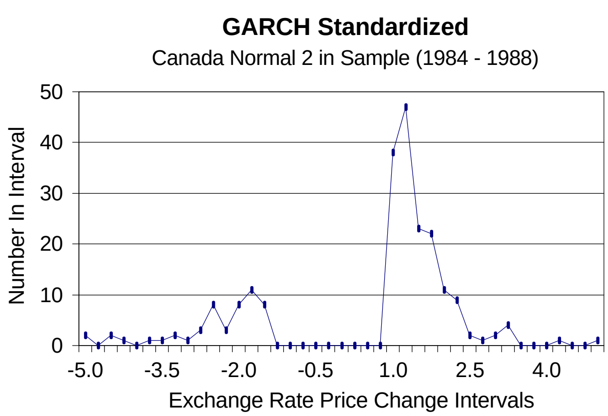

Sample Sizes: Overall = 1254 Normal 1 = 1046 Normal 2 = 208

Table 4, Panel A examines the two distributions for each standardization method, the null hypothesis is rejected both in and out of sample. This indicates that the data in each case does not come from the same distribution, and independence still exists in the data. These results show that independence needs to be incorporated in the stochastic process underlying exchange rate price changes and that GARCH type models alone are inadequate.

The Kolomogorov-Smirnov test is also used across standardization methods, Table 4, Panel B, to determine if the mean-variance standardized large distribution comes from the same distribution as the GARCH standardized large distribution. This is also done for the small distributions. Table 4 indicates that standardizing by different methods results in rejection of the null hypothesis for both the large and small distributions. Different standardization methods change the distributional form. The standardization methods result in significantly different distributions, therefore comparing these two series of observations will test the contribution of the conditional volatility since annual mean-variance standardization does not take into account conditional volatility.

TABLE 4
Non-Parametric Kolmogorov-Smirnov
Test For the Canadian Dollar
In Sample (1984-1988), Out of Sample (1985-1989)

Panel A: Ho: The 2 groups of observations have identical distributions.

	D	Critical Value of D (α = .01)
<u>Mean-Var Standardized</u>		
In Sample	.631578	.0460297
Out of Sample	.545455	.0458836
<u>GARCH Standardized</u>		
In Sample	.764423	.0460297
Out of Sample	.703704	.0458836

Panel B: Ho: Large (small) distributions based on the different standardization methods come from the same distributions.

	D	Critical Value of D (α = .01)
<u>In Sample</u>		
Large	.123097	.0327973
Small	.521002	.1039250
<u>Out of Sample</u>		
Large	.122060	.0343710
Small	.634613	.0937780

D is the test statistic for the non-parametric Kolmogorov-Smirnov test. It is given by: $D = Max_t|F_1(x_t) - F_2(x_t)|$ where $F_1(x_t)$ is the observed cumulative relative frequency for the t^{th} data value from distribution 1 and $F_2(x_t)$ is the observed cumulative frequency for the t^{th} data value from distribution 2. If D is greater than its critical value, the null hypothesis is rejected, thus supporting independence.

Table 5, Panel A presents the maximum likelihood estimation values, and the likelihood ratio tests comparing the GARCH standardized data with the mean-variance standardization data. In sample, both methods reduce the effect of outliers, one normal fits the data best using the Schwarz criteria (SC) to select the best fitting model. The Schwarz criterion (SC) controls for the degrees of freedom by giving more weight to models with fewer parameters.

Equation 7

$SC = log[L(x \quad \Theta)] - klog(\sqrt{N})$

Where $L(x \quad \Theta)$ is the value of the model's maximized likelihood function, N is the sample size and k is the number of independent parameters in the model.

For the raw data, consistent with previous research two normals is found to have a superior fit (not shown). Out of sample, mean-variance standardized data, two normals fits better, while one normal is again the best fitting model for the GARCH data.

Although there may be some increased reduction in the influence of outliers using the GARCH model, one normal always fits best, the likelihood ratio tests (Table 5, Panel B) indicate that mean-variance standardization is superior to GARCH standardization. The null hypothesis that the mean-variance standardization is superior to GARCH is accepted in all cases both in and out of sample and for one and two normals. Given that the GARCH coefficients are significant, if volatility clustering plays a significant role, we would expect the maximum likelihood value of the GARCH standardized data to be superior to the maximum likelihood value of the mean-variance standardized data. These findings show that mean-variance standardization is superior to using the GARCH model, again our findings put into question the contribution of conditional variance in explaining exchange rate price changes.

TABLE 5
Maximum likelihood Results and
Likelihood Ratio Tests, Canadian Dollar
In Sample (1984-1988) Out of Sample (1985-1989)

Panel A: Maximum Likelihood

	In Sample		Out of Sample	
	1 Normal	2 Normals	1 Normal	2 Normals
GARCH Standardized	1603.13	1491.29	1652.07	1547.15
(SC)	(1588.86)	(1477.02)	(1637.79)	(1532.87)
Mean-Var Standardized	1695.90	1536.47	1725.04	1814.75
(SC)	(1688.77)	(1518.63)	(1717.90)	(1796.90)

Panel B: Likelihood Ratio Tests

	In Sample	Out of Sample
	1 Normal	2 Normals
In Sample	185.54*	90.36*
Out of Sample	145.94*	535.20*

$SC = log[L(x \quad \Theta)] - klog(\sqrt{N})$

*Significant at the .01 level

The likelihood ratio test is $2(L_{mv} - L_g)$, where L_g is the GARCH standardized maximum log likelihood estimates and L_{mv} is the mean-variance standardized maximum log likelihood estimates. This test has a X^2 distribution with k degrees of freedom where k is the number of restrictions being tested.

CONCLUSION

In this study we examine the extent of the original GARCH models contribution to our understanding of the stochastic process underlying exchange rate changes, and determine if the movement of research to GARCH type models exclusively is warranted.

Overall results demonstrate that, although previous research indicates that volatility clustering plays a role in determining exchange rate price changes, it is not the primary factor generating these changes. GARCH models with normality assumptions do not provide a good description of exchange rate dynamics. Frequency distributions show independence still exists in the data after removing the GARCH effects. Likelihood ratio tests indicate that a simple alternative of using mean and variances to standardize the data is superior, to using the GARCH model. This study questions the contribution of GARCH type models, in the determination of the stochastic process underlying exchange rate price changes.

Results indicate that in some years when both GARCH conditional variance parameters are significant, tests for first-order heteroskedasticity are insignificant. A question for future research is what are the GARCH model $e^2(t\text{-}1)$ parameters explaining if there is no first-order heteroskedasticity?

This puzzle needs to be solved. When we are confident that the original GARCH model is explaining what it is theoretically designed to, we can move to more complex models.

Future research can examine if other forms of the GARCH process can account for the independence found (i.e., PGARCH, GARCH, FIGARCH, GARCH). They should also be tested to determine if they are superior to mean-variance standardization. Since all forms of the GARCH process are similar in form, focusing on volatility clustering, it would be interesting to see if they are an improvement.

ENDNOTES

1. The integrated GARCH model resembles ARIMA models for conditional first moments when obtaining optimal forecasts of the process.

2. Similar results are found with the other exchange rates.

3. Specification tests in the GARCH literature often are based on the residuals. Since the mean of the GARCH model is constant the use of the actual data in place of the residual would not affect results.

4. Out of sample frequency distributions are almost identical and therefore not shown.

REFERENCES

1. Baillie, R., Bollerslev, T., and Mikkelson, H., “Fractionally Integrated Generalized Autoregressive Conditional Heteroskedasticity,” *Journal of Econometrics*, September 1996, pp. 3-30.

2. Bollerslev, T., Chou, R., and Kroner, K., “Arch Modeling in Finance: A Review of the Theory and Empirical Evidence,” *Journal of Econometrics*, April-May 1992, pp. 5-59.

3. Bollerslev, T., and Ghysels, E., “Periodic Autoregressive Conditional Heteroskedasticity,” *Journal of Business and Economic Statistics*, April 1996, pp. 139-151.

4. Fujihara, R., and Park, K., “The Probability Distribution Of Futures Prices In The Foreign Exchange Market: A Comparison Of Candidate Processes,” *Journal of Futures Markets*, 1990, pp. 623-641.

5. Hsieh, D., “The Statistical Properties Of Daily Exchange Rates: 1974-1983,” *Journal of International Economics*, 1988, pp. 129-145.

6. Johnston, K., and Scott, E., “The Statistical Distribution of Foreign Exchange Rates: Dependant vs Independent Models,” *Journal of Financial and Strategic Decisions*, Vol. 12, No. 2, 1999, pp. 39-49.

7. Kugler, P., and Lenz, C., “Chaos, ARCH and The Foreign Exchange Market: Empirical Results From Weekly Data,” Unpublished Manuscript, Volkswirtschaftliches Institut, Zurich, 1990.

8. McCurdy, T., and Morgan, I., “Testing The Martingale Hypothesis In Deutsche Mark Futures With Models Specifying The Form Of The Heteroskedasticity,” *Journal of Applied Econometrics* 3, 1988, pp. 187-202.

9. Taylor, S., “Modeling Financial Time Series,” Wiley, New York, NY, 1986.