NORTH-HOLLAND

Financial Services Review 9 (2000) 219–230

FINANCIAL SERVICES REVIEW

Exploitable patterns in retirement annuity returns: evidence from TIAA/CREF

Edward M. Miller[a], Larry J. Prather[b,*]

[a]*Department of Economics and Finance, University of New Orleans, New Orleans, LA 70148, USA*
[b]*Department of Economics and Finance, East Tennessee State University, Box 70686, Johnson City, TN 37614, USA*

Received 18 September 2000; received in revised form 19 December 2000; accepted 7 February 2001

Abstract

Evidence suggests that predictabilities in asset class returns exist but transactions costs prevent exploiting them using individual securities. Extant research also shows that these relationships may by exploitable through the trading of mutual funds but fails to examine whether this relationship exists within an individual fund family. This paper finds that TIAA/CREF retirement annuities exhibit predictable elements that could be exploited by informed traders. The proposed trading strategy dominates a buy-and-hold strategy by producing higher raw and risk-adjusted returns. Additionally, the risk is greatly reduced. © 2001 Elsevier Science Inc. All rights reserved.

JEL classification: G12

Keywords: Return predictability; Technical trading rules; Lead and lag relationships among asset classes; Retirement annuities; Technical trading rules

1. Introduction

There is now substantial evidence that stock prices exhibit an appreciable amount of predictability (e.g., Chalmers, Edelen, & Kadlec, 1999; Goetzmann, Ivković, & Rouwenhorst, 2000; Hamao, Masulis & Ng, 1990; Lo & MacKinlay, 1999). This is in spite of the traditional efficient market theory argument that if an appreciable predictive element were

* Corresponding author. Tel.: +1-423-439-5668; fax: +1-423-439-8583.
E-mail addresses: Emmef@uno.edu (E.M. Miller); prather@etsu.edu (L.J. Prather).

PII: S1057-0810(01)00068-3

found, investors would trade so as to exploit the predictability. Miller and Prather (1999) find similar return predictabilities using mutual fund investment objectives as asset class proxies. However, extant research on asset class return predictabilities does not address any specific fund family. Rather, it explores the aggregate return predictabilities using asset class return indexes. Since Najand and Prather (1999) find that common fund classification schemes fail to capture all elements of risk, it is not possible to determine whether return predictabilities exist or whether a predictable component is exploitable within any given fund family.

This article addresses the issue of whether these predictabilities exist in TIAA/CREF retirement annuities. This line of investigation is important for several reasons. First, TIAA/CREF is one of the largest providers of retirement annuities to members of universities. Therefore, it is important to determine whether the effects found in theoretical portfolios constructed by researchers are observable in portfolios that fund managers construct and investors purchase. Secondly, if return predictabilities exist, can knowledgeable investors use these predictabilities to garner higher risk-adjusted returns through frequent trading? Evidence suggests that some investors may be attempting to move funds on a systematic basis. A recent article in *The Participant*, the periodic newsletter for TIAA/CREF participants, expressed concern that some investors are making frequent trades and that this is problematic for TIAA/CREF. Additionally, a new section was added to the College Retirement Equities Fund Prospectus, dated May 1, 2000, concerning the implementation of a market timing policy. This policy restricts the number of trades permitted in an effort to reduce the negative impact transactions costs incurred by TIAA/CREF due to this frequent trading. Since this policy would not prevent trading on the turn-of-the-month effect found by Kunkel and Compton (1998), it is possible another driver of frequent trading exists.

Empirical testing uses the Granger causality method to examine daily returns of retirement annuities during a 1453-day period. The Granger causality method enables us to test directly whether knowing past returns on one asset class will aid in the forecast of returns for another asset class. Next, we develop a trading strategy that attempts to exploit observed statistical relationships. Then we test the trading strategy on a holdout sample and find that it dominates a buy-and-hold strategy in terms of raw and risk-adjusted returns. Finally, we show that the new market timing policy implemented by TIAA/CREF does not eliminate the ability of informed traders to exploit predictable components of returns. As an example of the effectiveness of this trading rule, an initial investment of $100,000 at the beginning of the holdout period would have increased to $186,450 if invested entirely in the index fund whereas it would have grown to $253,180 if moved systematically. Thus, frequent trading would lead to a profit of $66,730 while reducing risk.

2. Background

While individual stocks show only a low predictability, when these stocks are combined into portfolios, the predictability may be higher or even different in nature. Lo and MacKinlay (1999) found this for small stocks. Small stocks had a low correlation with a lagged index of larger stocks and tended to have a negative autocorrelation over short periods. This is probably the result of idiosyncratic risk. However, when the stocks were combined into portfolios, they showed an appreciable positive correlation with the earlier performance of large stock indices. This results because the noise in the time series of individual stocks is diversified away when the stocks are combined into portfolios, leaving systematic effects to be revealed.

Miller and Prather (1999) show that studying mutual funds is one way to quickly study portfolios, since each fund represents a portfolio. Further, they show that, on average, exploitable regularities exist. This is important since mutual funds can be frequently traded with negligible costs. Within a fund family, transfers can usually be made by phone at no cost. Even between fund families, trades can often be made at zero or low cost. Usually one can sell shares in a mutual fund back to the fund for net asset value and the proceeds can then be reinvested in a new no-load fund at no expense. Retirement annuities function in much the same way as mutual funds; therefore, they are expected to exhibit similar patterns. A benefit of using retirement annuities to exploit predictable patterns in security returns, over individual stocks or even mutual funds, is the lack of taxes and transaction costs.

Trading mutual funds, or retirement annuities, lacks the same self-correcting forces that stock trading has. An example can illustrate the self-correcting nature of trading. Suppose that a rule is found that predicts that growth stocks will rise by an appreciable amount from Tuesday's close to Wednesday's close. A single investor may be able to profit from this rule by buying at the close on Tuesday and selling at Wednesday close. However, if many investors discover the rule, and buy at Tuesday's close and sell at Wednesday's close, they will bid the Tuesday's closing price up and lower Wednesday's closing price, tending to eliminate the effect.

The same effect that produces the predictability in the stock prices would be expected to produce predictability in the net asset values of a mutual fund or retirement annuity. Thus, an investor might be able to make a profit by buying the mutual fund, or retirement annuity, at the close on Tuesday (done by a phone call before the close) and then selling it at the close on Wednesday. Clearly, a single investor could profit from such predictability. However, large numbers of investors might also profit from it. Imagine investors moving a sum of money that would eliminate the effect if the trading were done directly in stocks. The same trading would be unlikely to eliminate the effect if the trading was done in mutual funds or retirement annuities.

The reason is that the fund manager does not have the influx of funds until after the close Tuesday (when the trades are actually done). Therefore, if the fund manager takes no action until the new funds are received, the effect could continue since no action was taken to affect stock prices. Alternately, suppose the fund manager responds to the influx of new funds by buying more stocks immediately. Two possibilities exist. First, if his buying is too little to affect the prices, the effect continues. However, if his buying is large enough to affect prices, his buying raises the prices of the stocks his fund buys, and the net asset value is even greater when it is calculated at Wednesday's close. The trading actually accentuated the price change (this would be especially likely for a specialized fund whose manager buys and sells within a small list of stocks that may be limited in liquidity, such as a country fund). Thus, day-to-day predictabilities could persist even if they were known.

3. Data and methodology

3.1. Data

In order to ascertain if predictabilities in retirement returns exist, we carry out empirical analysis. The sample consists of 1,453 daily returns on the following TIAA/CREF retirement annuities: Money Market, Bond, Stock, Index, Growth, and Global. Daily return data for each of the selected retirement annuities during the period May 2, 1994 through December 31, 1999 was obtained from TIAA/CREF. The objectives of these funds remained constant during the period of our study.

3.2. Computation of returns

Daily returns are computed for each retirement annuity by taking the change in net asset value (NAV) for an investment for each of the 1,453 days in our sample, as shown in Eq. (1)

$$R_{i,t} = \frac{NAV_{i,t} - NAV_{i,t-1}}{NAV_{i,t-1}} \tag{1}$$

where $R_{i,t}$ is the return on retirement annuity i during day t, $NAV_{i,t}$ is the net asset value of an investment in retirement annuity i at time t.

3.3. Methodologies employed

Several methodologies have been employed in the literature to ascertain lead and lag relationships. However, one of the most popular is the Granger causality test. The popularity of the Granger causality test stems from its ability to use an observable time series to test whether using information contained in that time series would be beneficial in forecasting another time series. Granger's (1969) is the only established method that permits testing whether one portfolio's returns are predictable by another portfolio's returns after controlling for autocorrelation. Therefore, it is useful in inferring the relative predictability between stochastic variables to ascertain a lead-lag structure. The Granger approach to the question of whether X causes Y is to determine the amount of the current Y that can be explained by past values of Y and then to ascertain whether adding lagged values of X can improve the explanation. Y is said to be Granger-caused by X if X helps in the prediction of Y, or equivalently if the coefficients on the lagged Xs are statistically significant. It is important to note that the statement "X Granger causes Y" does not imply that Y is the effect or the result of X. Granger causality measures information content but does not indicate causality in the common use of the term. Formally, a time series $\{X_t\}$ "Granger causes" another time series $\{Y_t\}$ if using past values of X can improve the forecast of Y. Since the objective of this research is to determine whether past returns from one retirement annuity are useful in predicting the future returns of another retirement annuity, this paper follows Richardson and Peterson (1999) and uses techniques based upon Granger causality to ascertain lead and lag relationships in returns between various retirement annuities.

Our test for return predictability uses Eqs. (2) and (3):

$$R_{i,t} = \alpha + \sum_{k=1}^{n} \beta_{i,k} R_{i,t-k} + \sum_{k=1}^{n} \beta_{j,k} R_{j,t-k} + \varepsilon_t \tag{2}$$

$$R_{j,t} = \delta + \sum_{k=1}^{n} \gamma_{i,k} R_{i,t-k} + \sum_{k=1}^{n} \gamma_{i,k} R_{j,t-k} + \upsilon_t \tag{3}$$

where n is the number of lags estimated; $R_{i,t}$ is the return series for asset class i; $R_{j,t}$ is the return series for asset class j; α and δ are the estimated intercepts; $\beta_{i,k}$ and $\gamma_{i,k}$ are the coefficients for asset class i's return series lagged t-k periods; $\beta_{j,k}$ and $\gamma_{j,k}$ are the coefficients for asset class j's return series lagged t-k periods; and ϵ_t and υ_t are the normally distributed error terms.

The Granger causality tests consist of whether all the coefficients of the lagged Xs in Eq. (2) may be considered to be zero, and similarly whether the coefficients of the lagged Ys in Eq. (3) are zero. Thus, the null hypotheses being tested are that X does not Granger-cause Y and that Y does not Granger-cause X.

4. Empirical results of lead and lag relationships among asset classes

In order to conduct empirical investigation, we first divided our sample into two subsamples with an approximately equal number of observations. This division allows testing lead and lag relationships and the development of trading rules with one subsample and then testing the dominance of those rules over a buy and hold strategy using the holdout sample. The first subsample contains 725 daily observations from May 2, 1994 through February 28, 1997. The holdout sample contains 728 daily observations from March 3, 1997 through December 31, 1999.

The first step in assessing whether it is possible to utilize return patterns from one asset class to foretell the average future return of another asset class is to determine the instantaneous correlation. If instantaneous correlations are high, the impact of any trading strategy is mitigated.

Table 1 provides the instantaneous correlation and covariance matrices for the six asset classes in our sample. The instantaneous correlation table is of interest in its own right since it shows that diversification between different asset classes is possible. The relatively low correlations between Global funds and other categories of funds suggest that Global funds can indeed be used to reduce risk when combined with domestic U.S. funds. The highest correlation (0.974) is between Index and the Stock funds. The high correlation between the index and the stock fund is because approximately two thirds of the stock fund is indexed.

The instantaneous correlation structure suggests that investors may be able to benefit from an asset reallocation strategy if observing the returns on one asset class could provide information about the future returns of another asset class.

Table 1
Instantaneous correlation and covariance between asset classes

	Panel A Instantaneous correlation					
	STK	MM	BOND	GLOB	GRO	INDX
STK	1.000000	0.036410	0.530942	0.742750	0.942812	0.973581
MM	0.036410	1.000000	0.110867	0.007747	0.004376	0.030099
BOND	0.530942	0.110867	1.000000	0.269797	0.474017	0.551139
GLOB	0.742750	0.007747	0.269797	1.000000	0.629086	0.607968
GRO	0.942812	0.004376	0.474017	0.629086	1.000000	0.961849
INDX	0.973581	0.030099	0.551139	0.607958	0.961849	1.000000

	Panel B Instantaneous covariance matrix					
	BOND	GLOB	GRO	INDX	MM	STK
BOND	7.42E-06	3.63E-06	8.67E-06	9.11E-06	3.81E-08	7.41E-06
GLOB	3.63E-06	2.44E-05	2.09E-05	1.82E-05	4.82E-09	1.88E-05
GRO	8.67E-06	2.09E-05	4.51E-05	3.92E-05	3.70E-09	3.24E-05
INDX	9.11E-06	1.82E-05	3.92E-05	3.68E-05	2.30E-08	3.03E-05
MM	3.81E-08	4.82E-09	3.70E-09	2.30E-08	1.59E-08	2.35E-08
STK	7.41E-06	1.88E-05	3.24E-05	3.03E-05	2.35E-08	2.62E-05

4.1. Tests of Granger causality

Since investors may be able to benefit from an asset reallocation strategy if observing the returns on one retirement annuity could provide information about the future returns of another retirement annuity, we examine whether that situation exists (on average) using Granger causality. Table 2 provides the results of Granger causality tests of lead-lag

Table 2
Pairwise Granger causality tests for a one-period lag

	STK	MM	BOND	GLOB	GRO	INDX
STK		5.2915	0.4090	60.4311	0.6969	0.2362
		(0.0217)	(0.5227)	(0.0000)	(0.4041)	(0.6272)
MM	0.6290		3.6568	0.2450	0.9709	0.6189
	(0.4280)		(0.0563)	(0.6208)	(0.3248)	(0.4317)
BOND	0.4412	4.4005		32.5721	0.0833	0.1137
	(0.5068)	(0.0363)		(0.0000)	(0.7730)	(0.7361)
GLOB	0.0474	2.5997	3.0471		0.7126	1.0155
	(0.8277)	(0.1073)	(0.0813)		(0.3989)	(0.3139)
GRO	0.0390	2.9648	0.0051	59.4262		0.1944
	(0.8435)	(0.0855)	(0.9431)	(0.0000)		(0.6594)
INDX	1.4763	4.7076	0.0929	69.0951	1.2504	
	(0.2247)	(0.0304)	(0.7607)	(0.0000)	(0.2639)	

A matrix of the F-statistic (p-value in parenthesis) for the Granger causality test of lead and lag relationships is presented for the initial sample of 725 trading days between May 2, 1994 and February 28, 1997. The test statistics are for testing the null hypothesis that the asset class returns in the row do not Granger cause the asset class returns in the column for a one-day lag.

relationships of the thirty possible test pair combinations over the 725-day sample period. These tests of Granger causality are stated in null hypothesis form and examine a one-period (trading-day) lag. Interestingly, seven test pairs yielded significant lead-lag relationships at the five-percent level (the null hypothesis was rejected) and three additional test pairs were significant at the ten-percent level. The finding of seven significant causalities at the five-percent level is far above the one or two that would be expected from chance alone. This suggests that trading patterns may exist and that investors may develop a personal strategy to maximize their utility by using this lead-lag information to enhance the movement of funds among preselected retirement annuities that provide the desired level of risk.

One caveat exists. Granger causing a relationship does not mean one asset class causes another, it simply means that a statistical relationship exists such that knowing past returns of one class will help determine future returns in another. Empirical analysis suggests that the most significant statistical relationships are for: (1) Index leading Global, (2) Stock leading Global, (3) Growth leading Global, (4) Bond leading Global, (5) Stock leading Money Market, (6) Index leading Money Market, and (7) Bond leading Money Market. Each of these relationships is significant at better than the five-percent level.

4.2. Practical limitations

As Miller and Prather (1999) point out, one limitation of trading on a one-day lag is that it presupposes that an investor could determine returns on the asset class during any given day, sell the fund at the close, and immediately reposition the assets. However, since this study is using retirement annuities as the securities in the asset class instead of individual stocks, this is not possible. The returns are not known with certainty until after the close of business. This permits the fund to compute net asset value (NAV) after closing prices of the securities in the portfolio are determined. Sell orders are executed at the closing NAV on the day the order is processed and buy orders are executed at the closing NAV on the day the buy order is processed. Therefore, if after observing past changes in NAV, an order were placed to sell/buy, the prices would be the prices on event day $t + 1$, not day t.

Miller and Prather (1999) deal with this complication by examining relationships based on a two-day lag and form a conservative trading strategy using the two-day lagged returns. However, they state that it is reasonable to believe that investors could monitor index returns and make decisions to trade based on the index that proxies for their current asset class. We find a similar two-day lag relationship with TIAA/CREF data (not reported in the paper). However, we base our trading rule on the ability to monitor the index and therefore, use a one-day lag to report results. This is explained in the following section.

4.3. Exploitation of return predictability

Exploiting information provided by the Granger causality tests requires moving assets based on the strength of the statistical relationship. Therefore, it is necessary to examine the Granger F-statistic (p-value) to find the strongest statistical relationship. Results in Table 2 suggest that the strongest relationship is that Index causes Global. This is consistent with Hamao, Masulis and Ng (1990) who document spillover effects between U.S. and international markets. One explanation for this is that it arises from an asynchronous pricing problem (e.g., Chalmers, Edelen, & Kadlec, 1999; Goetzmann, Ivković, and Rouwenhorst, 2000; Varela, 1997). When one buys a fund the net asset value is based on the latest available prices as of the time the New York stock exchange closes (4 p.m. Eastern time). However, for Asian and European stocks the latest available prices are many hours old. Some pieces of news that have affected the U.S. stocks prices probably have also affected the values of these foreign stocks, but the prices will not actually change until the Asian and European markets open the next day. This creates an opportunity to trade at stale prices.

This relationship is fortunate for traders because the returns on an index fund closely approximate the returns on the index it tracks. The TIAA/CREF prospectus (May 2000) indicates that the Index Fund tracks the Russell 3000 index. The Russell 3000 index is an unmanaged index of 3,000 of the largest U.S. companies (based on market capitalization). The Index Fund does not hold all 3000 stocks but uses a sampling technique to closely replicate the index. The goal of TIAA/CREF management is to attempt to match the total return on the Russell 3000 Index but with any index fund, they may not always do so. As a practical matter, the tracking error ($1\text{-}R^2$) of index funds is typically very small. While the values of the Index Fund are calculated only at the end of the trading day, the index itself is available during the day. Thus, the result of a strategy of making decisions on the closing price of the Index Fund can be approximated by checking the index just before the closing, and then trading based on the performance of the index. This is important since TIAA/CREF permits telephone or the Inter/ACT Internet service transactions. Any transactions received before the market closes will be executed that day.

The TIAA/CREF prospectus indicates that the Global Fund invests at least 65% of its assets in equity securities of foreign and domestic companies. Typically, at least 40% of assets are allocated to foreign securities and 25% of assets are allocated to domestic securities. The remaining 35% of assets are distributed between foreign and domestic securities based upon market conditions. The CREF Global fund usually has only a minority of its funds in foreign stocks. Many other fund families include foreign stock funds that hold most of their equity positions in Asian and European stocks. If there is an exploitable effect for the CREF funds, a similar strategy would probably be even more profitable for many other fund families and retirement plans that permit frequent trading.

Since investors can monitor the Russell 3000, our trading strategy is based on a one-day lag. This strategy calls for moving assets from Index to Global on a positive return in Index. This move is made since the investor can use knowledge of return predictability to predict a positive Global return. The investor stays in Global until a negative return is observed for Index. Upon observing a negative return on Index, the investor returns to Index (to avoid a negative return in Global during the upcoming period). Since Global does not lead Index, there is no reason to believe that Index will be negative in the next period. This pattern is consistent with the findings of Hamao, Masulis and Ng (1990).

It is worthwhile examining the magnitude of the predictable component to determine the potential gain from trading. Table 3 provides results of regressing returns of the Global annuity on one-day lagged returns of the Index annuity. Results suggest that a positive one-percent return by the Index annuity on day t should lead to an average positive return on the Global annuity of 0.3% return on day $t + 1$.

Table 3
Predictable return components of selected trading strategy

INDEX leading GLOBAL					
Variable	Coefficient	Std. Error	T-Statistic	Prob.	R^2
INDEX(-1)	0.302026	0.028150	10.72927	0.0000	0.1375
α	0.000286	0.000172	1.65577	0.0982	

The regression of Global returns on the one-day lagged predictor variable (Index) is presented to determine the magnitude of the predictable element of returns. Columns one through six present the variable, coefficient, standard error, T-statistic, probability (p-value), and coefficient of determination (R^2) of the regressions, respectively.

A practical constraint faced by an informed trader is that TIAA/CREF will not permit frequent movement of assets from one asset class to another. The CREF prospectus dated May 1, 2000, states "participants who make more than three transfers out of any account (other than the Money Market Account) in a calendar month will be advised that if this transfer frequency continues, we will suspend their ability to make telephone, fax, and internet transfers" (p. 36). This places constraints on knowledgeable investors attempting to capitalize on return predictabilities. An examination of the data shows that the Index returns were positive on 387 days in the second sample period (53%). Returns for Global were similar with 414 positive daily returns (57%). Additionally, Global had a positive return on 254 of the 387 trading days following a positive return on the Index (66%). Therefore, an examination of the return distributions is required prior to attempting to form any implementable rule since the frequent trading suggested by this simple rule would be prohibited by TIAA/CREF.

In order to limit the number of trades, the sample was examined to ascertain the characteristics of the distribution of returns. The goal is to be able to trade only on the most positive Index returns, not just any positive Index return. As a starting point, we selected a return slightly higher than the top twenty percent (0.0093) as the trigger point for moving into Global. Prior to executing the strategy, we also examined the number of returns in this top twenty percentage category and how frequently positive Index returns were followed by positive Global returns. Of the series of returns, 146 were above the selected trigger point. Interestingly, 99 of the 146 large positive Index returns were followed by positive Global returns (68%). We also selected a trigger point to move back into Index based on negative Index return. Again, the trigger point was a bottom twenty-percent return (−0.0068). This type of constrained trading should be useful in reducing the number of trades and avoiding large negative returns, while capitalizing on expected large positive returns.

4.4. Risk and return of trading strategies

The returns, risks, and Sharpe (1966) measures for both buy-and-hold strategies and the proposed trading rule are presented in Table 4. The Sharpe measure (S) is computed as $S = [Rp\text{-}R_f]/\sigma_p$. The trading rule portfolio has the highest average daily return of any of the portfolios (0.00142) compared to 0.00086 for Stock (lowest equity return) and 0.00113 for

Table 4
Risk and return of sample portfolios

Asset Class	Return	Std. Dev.	Sharpe
STK	0.00086	0.01015	0.06535
BOND	0.00022	0.00230	0.00732
GLOB	0.00091	0.00921	0.07639
GRO	0.00113	0.01298	0.07188
INDX	0.00092	0.01132	0.06346
RULE	0.00142	0.01062	0.11488

Column one is the asset class, columns two and three are the average arithmetic return and standard deviation of returns. Column four provides the Sharpe measure.

Growth (highest equity return). It also has a lower total risk, as measured by standard deviation (0.01062), than either the Index (0.01132) or the Growth portfolios (0.01298). This characteristic leads to the Sharpe measure of the trading rule portfolio (0.11488) being higher than any buy-and-hold strategy. None of these outcomes are surprising given the strong Granger relationship. Trading on Granger relationships has increased the probability of obtaining a positive return in the period following a positive return in the Index and reduced the probability of receiving a large negative return in Global. This tends to increase average returns and to reduce variability of returns (risk), which leads directly to the higher Sharpe measure.

As a test of robustness, the Jensen (1968) measure of risk-adjusted returns was computed to determine whether the positive risk-adjusted returns were statistically significant. The Jensen measure is computed as $R_p\text{-}R_f = \alpha + \beta_p (R_m - R_f) + \epsilon$. For our purposes, the ordinary least squares regression uses the CREF Money Market returns as the risk-free proxy and the CREF Index returns as the market proxy. This approach was selected since TIAA/CREF investors can form a portfolio using these asset classes to achieve their desired level of systematic risk. This is a realistic approach since it represents actual returns investors could achieve after costs. Table 5 shows that the model fit is good since more than 88% of the trading rule returns are explained by the model. This is comforting since Brown and Brown (1987) and Lehmann and Modest (1987) find that the market proxy is important when attempting to determine true performance. Risk-adjusted returns of the trading rule are positive and statistically significant at better than the one-percent level. Additionally, systematic risk of the trading rule ($\beta = 0.88$) is less than that of the Index.

Table 5
Risk-adjusted return of trading rule

α	t-statistic	p-value	β	R^2	N
0.00059 (.00013)	4.40436	0.00001	0.88287 (0.01173)	0.88637	728

The results of the market model regression are presented. Columns one through three present the risk-adjusted return (α), the t-statistic, and p-value for the two-tailed hypothesis test that risk-adjusted return equals zero. Columns four through six present the systematic risk (β), coefficient of determination (R^2), and number of observations (N). Standard errors are in parenthesis below the coefficient estimates.

Given the superior Sharpe and Jensen performance measures for the trading strategy, it appears that informed investors were able to exploit asset class return predictabilities in TIAA/CREF family retirement annuities. To achieve this result, 139 trades were conducted over 728 trading days for an average trading frequency of approximately 3.8 trades per month. Given the current restrictions that limit trading out of a given annuity to three trades per month, it is important to determine if this constraint had been violated by our trading rule.

Examination of the trading pattern reveals that in seven of the 34 months the trading limit would have been reached in at least one of the annuities and that in three months, some trades may have been prohibited. To ascertain the impact of this, we took a worst-case approach and assumed that the trades would have been prohibited. Therefore, the investor would be required to remain in the current asset class once the trading limit was reached. Results show that the limits had little impact on previous results. The average return for the constrained trading strategy is 0.00133 and the standard deviation is 0.010574, which produces a Sharpe measure of 0.106961. This compares favorably with the results of buy-and-hold strategies in Table 4 and is slightly less than the unconstrained rule. Results of the Jensen measure are also similar to the unconstrained rule. Risk-adjusted return is 0.0005 and statistically significant with a p-value of 0.00018. Additionally, risk and model fit are essentially unchanged from the unconstrained model.

5. Conclusion

Recent evidence suggests that asset class returns possess a predictable component. However, transactions costs such as commissions and bid-ask spreads prevent investors from exploiting this predictability.

This study examines asset class returns of TIAA/CREF retirement annuities for several reasons. First, investors can trade these retirement annuities without incurring trading costs. Additionally, these annuities should exhibit return predictabilities similar to mutual funds and recent evidence suggests that mutual fund asset class indices exhibit a predictable component. Our primary goal is to learn whether TIAA/CREF retirement annuities exhibit asset class return predictability and to see if a profitable trading scheme exists.

We performed Granger causality tests and found that the retirement annuity asset class returns exhibited a predictable element. Using the strength of the predictable element, we developed a trading strategy and tested that trading rule on a holdout sample.

Examination of risks and returns of the trading rule and buy-and-hold strategies show that the trading rule has superior Sharpe and Jensen performance measures compared with a buy-and-hold strategy. Further, the recent trading constraints implemented by TIAA/CREF do not eliminate the ability of informed investors to enhance their reward-to-risk ratio through trading based on statistical relationships. This finding has implications for management at TIAA/CREF, portfolio managers, and investors seeking the best risk-return relationship.

References

Brown, K., & Brown, G. (1987). Does the composition of the market portfolio really matter? *Journal of Portfolio Management,* (Winter), 26–32.

Chalmers, J., Edelen, J., & Kadlec, G. (1999). The Wildcard Option in Transacting Mutual-Fund Shares. Unpublished working paper, University of Pennsylvania.

College Retirement Equities Fund Prospectus, Individual, Group, and Tax-Deferred Variable Annuities. (May 1) 2000, pp. 1–45.

Goetzmann, W., Ivković, Z., & Rouwenhorst, G. (2000). Day trading international mutual funds: Evidence and policy solutions. Unpublished working paper, Yale University.

Granger, C. (1969). Investigating causal relationships by econometric models and cross-spectral models. *Econometrica, 37,* 424–438.

Hamao, Y., Masulis, R., & Ng, V. (1990). Correlations in price changes and volatility across international stock markets. *The Review of Financial Studies, 3,* 281–307.

Jensen, M. (1968). The performance of mutual funds in the period 1945–1964. *Journal of Finance, 23,* 389–419.

Kunkel, R., & Compton, W. (1998). A tax-free exploitation of the turn-of-the-month effect: C.R.E.F. *Financial Services Review, 7,* 11–23.

Lehmann, B., & Modest, D. (1987). Mutual fund performance evaluation: a comparison of benchmarks and benchmark comparisons. *Journal of Finance, 42,* 233–265.

Lo, A., & MacKinlay, A. (1999). *A non-random walk down Wall Street.* Princeton, New Jersey: Princeton.

Miller, E., & Prather, L. (1999). Lead and lag relationships among asset classes: evidence from the mutual fund industry. (Working Paper, East Tennessee State University).

Najand, M., & Prather, L. (1999). The risk level discriminatory power of mutual fund investment objectives: additional evidence. *Journal of Financial Markets, 2,* 307–328.

Sharpe, W. (1966). Mutual fund performance. *Journal of Business, 39,* 119–138.

Richardson, T., & Peterson, D. (1999). The cross-autocorrelation of sized-based portfolio returns is not an artifact of portfolio autocorrelation. *The Journal of Financial Research, 22,* 1–13.

Varela, O. (1997). Efficient market implications of institutional practices in obtaining net asset values for Asian-market based mutual funds in the U.S. Unpublished working paper No. 8–97, University of New Orleans).