Quarterly Journal of Finance
Vol. 1, No. 3 (2011) 465–493

DOI: 10.1142/S2010139211000122

Time-Varying Sharpe Ratios and Market Timing

Yi Tang
School of Business, Fordham University
Bronx, NY 10458, United States

Robert F. Whitelaw*
Stern School of Business
NYU, 44 W. 4th St., Suite 9-190
New York, NY 10012

National Bureau of Economic Research, USA

rwhitela@stern.nyu.edu

This paper documents predictable time-variation in stock market Sharpe ratios. Predetermined financial variables are used to estimate both the conditional mean and volatility of equity returns, and these moments are combined to estimate the conditional Sharpe ratio, or the Sharpe ratio is estimated directly as a linear function of these same variables. In sample, estimated conditional Sharpe ratios show substantial time-variation that coincides with the phases of the business cycle. Generally, Sharpe ratios are low at the peak of the cycle and high at the trough. In an out-of-sample analysis, using 10-year rolling regressions, relatively naive market-timing strategies that exploit this predictability can identify periods with Sharpe ratios more than 45% larger than the full sample value. In spite of the well-known predictability of volatility and the more controversial forecastability of returns, it is the latter factor that accounts primarily for both the in-sample and out-of-sample results.

Keywords: Sharpe ratio; predictability; stock market.

1. Introduction

The empirical literature contains a wealth of evidence on predictable variation in the mean and volatility of equity returns.[1] Given the apparent joint

*Corresponding author.
[1] See, for example, Breen *et al.* (1989), Fama and French (1989), Kandel and Stambaugh (1990), Keim and Stambaugh (1986), and Schwert (1989).

predictability of the mean and volatility, it is perhaps somewhat surprising that the literature has been relatively silent on predictable variation in equity market Sharpe ratios.[2] Of course, predictable variation in the individual moments does not imply predictable variation in the Sharpe ratio. The key question is whether these moments move together, leading to Sharpe ratios which are more stable and potentially less predictable than the two components individually. The empirical evidence on this issue is somewhat mixed. Earlier work (e.g., French *et al.*, 1987) suggests a weak positive relation between expected returns and volatility. However, other studies (e.g., Glosten *et al.*, 1993; Whitelaw, 1994; Boudoukh *et al.*, 1997) appear to uncover a more complex relation. Specifically, on an unconditional basis, several empirical specifications, including a modified GARCH-M model and nonparametric kernel estimation, suggest a negative relation between the conditional mean and volatility of returns. This evidence would indicate the likelihood of substantial predictable variation in market Sharpe ratios.

Time-variation in stock market Sharpe ratios is of interest for a number of reasons. First, the Sharpe ratio is popular for performance evaluation in an asset management context, and the unconditional Sharpe ratio of the market is often used as a convenient benchmark. If this ratio shows substantial predictable variation, then this variation needs to be accounted for when using the market as a performance benchmark. Second, time-variation in the market Sharpe ratio might provide clues to the fundamental economics underlying the economy and asset pricing. For example, the Sharpe ratio could indicate the timing and magnitude of fluctuations in risk aversion in a representative agent framework. Third, mean-variance investors would have an obvious interest in predictable Sharpe ratios as this variation could potentially lead to optimal trading strategies that differ markedly from simple buy-and-hold strategies. For example, in a partial equilibrium setting, the Sharpe ratio determines the fraction of wealth that an agent invests in the risky market portfolio.[3]

This paper provides an empirical investigation of time-variation in monthly, equity market Sharpe ratios over the sample period May 1953 to December 2010. We employ two different, but related, methodologies to

---

[2] An exception is Kandel and Stambaugh (1990), who investigated the implications of conditional moments of consumption growth for the price of risk over one-quarter and five-year horizons.

[3] See, for example, Kandel and Stambaugh (1996), who focused on the predictability of expected returns and its effect on asset allocation, and Fleming *et al.* (2001), who focused on the value of information about time-varying variances and covariances of returns.

come up with four different estimated conditional Sharpe ratios. In the first methodology, the conditional mean and volatility of equity returns are modeled as linear functions of the four predetermined financial variables used in Whitelaw (1994) plus lagged realized volatility. The ratio of these moments provides the estimate of the conditional Sharpe ratio. In order to decompose the predictability of the Sharpe ratio, we also consider two special cases of the above estimate in which either the mean or volatility is fixed at its unconditional average, and the other moment is allowed to vary over time. The second methodology estimates the conditional Sharpe ratio directly by projecting the monthly realized Sharpe ratio on to the same set of five variables. We examine the forecasting power of these estimates on both an in-sample and out-of-sample basis.

In sample estimated conditional Sharpe ratios exhibit substantial time-variation, with monthly values generally in the range of -0.2 to 0.9. Estimates from the ratio of the conditional moments exhibit similar patterns to those from direct estimation of the Sharpe ratio, with a correlation between the two series of approximately 0.9, but the latter are substantially higher due to the negative correlation between returns and volatility. Interestingly, the fixed-volatility estimate tracks the unconstrained estimates closely, while the fixed-mean estimate is much less variable and less correlated with its counterparts that allow expected returns to vary over time. Thus, these unconstrained estimates are clearly driven primarily by variation in the conditional expected return. The unconstrained conditional Sharpe ratios vary with the business cycle, peaking at business cycle troughs and declining over the course of the expansionary phase of the cycle. For example, the average increase in the conditional Sharpe ratio, estimated as the ratio of the conditional moments of returns, between the peak of the cycle and the subsequent trough is 0.25 (more than 0.85 on an annualized basis); therefore, the estimated risk-return tradeoff is much more favorable after recessions. Regressions of realized Sharpe ratios on the estimates show substantial predictive power, but the conditional unbiasedness of the forecasts can generally be rejected.

On an out-of-sample basis, using 10-year rolling regressions, estimated conditional Sharpe ratios again show statistically and economically significant predictive power for realized Sharpe ratios in spite of the fact that the rolling coefficients exhibit substantial instability. However, the slope coefficients in regressions of the realized ratio on the forecasts are much lower than for the in-sample exercise and are much lower than one, indicating substantial overfitting. Nevertheless, the forecasts provide valuable information in the context of a relatively naive market-timing strategy, significantly outperforming a buy-and-hold strategy. These active trading strategies involve switching between the market and the risk-free asset depending on the level of the estimated Sharpe ratio relative to a specific threshold. The average realized Sharpe ratios of the months in the market are compared to that of the buy-and-hold strategy, and they exhibit improvements of as much as 45%. Using ex post Sharpe ratios computed using monthly returns, these improvements are even more impressive, ranging up to 150%. Of equal importance, as the threshold for the conditional Sharpe ratio increases, so does the Sharpe ratio of the active strategy. Interestingly, forecasts based on the ratio of the conditional moments appear to perform the best, and there is little difference between the forecast that models volatility with the full set of conditioning variables and the one that uses simply the average volatility over the prior 10 years. This result complements results showing the economic value of volatility timing in an asset allocation context (e.g., Fleming *et al.*, 2001).

There are two possible interpretations of these results. First, they could be a product of mispricing, perhaps induced by fluctuations in consumer sentiment associated with business cycle fluctuations. In other words, at the peak of the cycle, over-optimistic investors could be overpricing stocks leading to poor future tradeoffs between risk and return, with the reverse happening at the trough of the cycle. If so, this paper provides a framework for analyzing and exploiting this inefficiency in order to generate superior performance. Second, the empirical results could be due to rational, but not perfectly correlated, time-variation in the conditional moments of returns. For example, if aggregate risk aversion changes over the cycle, decreasing during expansions and increasing during recessions, as might be implied by a habit model such as Campbell and Cochrane (1999), that would produce results generally in line with those reported in the paper. Alternatively, the world may resemble the ICAPM framework of Merton (1973), where time-varying investment opportunities generate an additional source of priced risk. For example, Whitelaw (2000) develops an equilibrium model in which the mean and volatility of market returns do not move together, implying substantial, rational, time-variation in stock market Sharpe ratios due to time-varying probabilities of regime switches.

The remainder of the paper is organized as follows. Section 2 provides a theoretical discussion and a setting to interpret time-varying Sharpe ratios. Section 3 introduces the estimation methodology and documents the economic and statistical significance of in-sample time-variation in stock market

Sharpe ratios. In Section 4, the out-of-sample analysis is performed, and the performance of stylized market-timing strategies is examined. Section 5 concludes.

2. Theoretical Background

Harrison and Kreps (1979) show that the absence of arbitrage implies the existence of a pricing kernel, or stochastic discount factor (M_t), that prices all assets. Specifically, the expected value of the product of the pricing kernel and the gross asset return (R_t) must equal unity, i.e.,

$$E_t[M_{t+1}R_{t+1}] = 1, \tag{1}$$

where E_t is the expectation based on information available at time t. Applying Equation (1), the one-period risk-free interest rate (R_{ft}) can be written as the inverse of the expectation of the pricing kernel

$$R_{t+1} = E_t[M_{t+1}]^{-1}. \tag{2}$$

Equation (1) also implies that the expected risk premium on any asset is proportional to the conditional covariance of its return with the pricing kernel, i.e.,

$$E_t[R_{t+1} - R_{ft}] = -R_{ft}\,\mathrm{cov}_t[M_{t+1}R_{t+1}], \tag{3}$$

where cov_t is the covariance conditional on information available at time t. Consequently, the conditional Sharpe ratio of any asset, defined as the ratio of the conditional mean excess return to the conditional standard deviation of this return, can be written in terms of the volatility of the pricing kernel and the correlation between the pricing kernel and the return as shown in Equation (4).

$$\frac{E_t[R_{t+1} - R_{ft}]}{\sigma_t[R_{t+1} - R_{ft}]} = \frac{-R_{ft}\mathrm{cov}_t[M_{t+1}R_{t+1}]}{\sigma_t[R_{t+1}]} = -R_{ft}\sigma_t[M_{t+1}]\,\mathrm{corr}_t[M_{t+1}R_{t+1}], \tag{4}$$

where σ_t and corr_t are the standard deviation and correlation, conditional on information at time t, respectively.

Denoting the conditional Sharpe ratio of the stock market at time t by S_t, Equation (4) shows that this ratio is proportional to the product of the volatility of the pricing kernel and the correlation between the pricing kernel and the return on the market (R_{mt})

$$S_t \equiv \frac{E_t[R_{mt+1} - R_{ft}]}{\sigma_t[R_{mt+1} - R_{ft}]} = -R_{ft}\sigma_t[M_{t+1}]\,\mathrm{corr}_t[M_{t+1}R_{mt+1}], \tag{5}$$

Intuitively, if the Sharpe ratio varies substantially over time, then this variation is attributable to variation in the conditional volatility or conditional correlation. Note that R_{ft} is the gross, risk-free rate, which has varied between 1.00 and 1.02 for monthly, US data.

Consider first the conditional correlation in Equation (5). The implications for time-variation in the Sharpe ratio depend critically on the modeling of the pricing kernel. One approach is to specify M_t as a function of asset returns. For example, modeling the pricing kernel as a linear function of the market return produces the conditional CAPM. Risk aversion implies a negative coefficient on the market return; therefore, the correlation is -1 and the market Sharpe ratio is approximately constant over time. Alternatively, modeling the discount factor as a quadratic function of the market return gives the conditional three-moment CAPM, first proposed by Kraus and Litzenberger (1976).[4] This specification allows for some time-variation in market Sharpe ratios due to the pricing of skewness risk, but the correlation will still be pushed towards -1. Bansal and Viswanathan (1993) estimate the pricing kernel as a general, nonlinear function of the market return, but again time-variation in the correlation is limited by a specification which relies on variation in market returns to proxy for variation in the discount factor. A slightly different approach to generalizing the one-factor, conditional CAPM, without abandoning a linear specification, is to model the pricing kernel as a linear function of multiple asset returns. Based on explanatory and predictive power, a number of additional factors, including small firm returns and return spreads between long-term and short-term bonds, have been proposed and tested.[5] However, as above, correlations between the discount factor and the market return tend to be relatively stable, implying stability in the stock market Sharpe ratio.

A second branch of the literature, using results from a representative agent, exchange economy (Lucas, 1978), models the pricing kernel as the marginal rate of substitution over consumption. The resulting consumption CAPM has been analyzed and tested in numerous contexts.[6] While the consumption CAPM literature is voluminous, there has been relatively little

[4]Note that these models can also be derived by imposing restrictions on a representative agent's utility function. Harvey and Siddique (2000) provide an empirical investigation of these specifications and a more detailed discussion of the underlying theory.

[5]Among the numerous papers on this topic are Campbell (1987), Chen *et al.* (1986), and Fama and French (1993).

[6]See, for example, Hansen and Singleton (1983), Breeden *et al.* (1989), and Ferson and Harvey (1992).

attention paid to the implications of the model for the risk/return relation.[7] Intuitively, when the marginal rate of substitution depends on consumption growth and the stock market is modeled as a claim on aggregate consumption, one might expect the correlation and the Sharpe ratio to be relatively stable. In fact, Whitelaw (2000) shows that this result holds when consumption growth follows an autoregressive process. However, in a two-regime model, where mean consumption growth differs across the regimes and the probabilities of regime shifts are time-varying, this intuition is overturned. In this setting, the mean and volatility of market returns can be negatively correlated. Although the magnitude and variation of the market Sharpe ratio are not investigated, it is clear that the model implies economically significant time-variation. It is important to note that the regimes correspond loosely to the expansionary and contractionary phases of the business cycle. Moreover, volatility and expected returns at any point in time depend critically on the probability of a regime shift. Consequently, this model predicts business cycle related variation in Sharpe ratios and large movements around transitions between the phases of the cycle.

A different approach is to modify the preferences of the representative agent. In the external habit model of Campbell and Cochrane (1999), utility depends on the deviation of consumption from a reference level. As consumption falls towards this reference level, for example during a recession, the effective risk aversion of the agent can increase dramatically, thus increasing the volatility of the pricing kernel in Equation (5). Through this type of mechanism, time-varying risk aversion can create business cycle variation in the Sharpe ratio.

3. Time-Variation in Market Sharpe Ratios

In this section, we first provide a brief description of the data, and we then turn to an explanation of the methodologies for estimating the conditional Sharpe ratio of market returns. The following subsection presents the in-sample empirical results.

3.1. *The data*

For this analysis, both the mean and volatility of stock market returns are estimated as functions of predetermined financial variables. Four

[7] One exception is Kandel and Stambaugh (1990), who constructed a four-state model of the mean and volatility of consumption growth. In their framework, the price of risk shows business cycle variation at long horizons due to variation in investment opportunities.

variables — the Baa-Aaa spread (DEF), the commercial paper-Treasury spread (CP), the one-year Treasury yield (1YR), and the dividend yield (DIV) — are chosen based on their proven predictive power in earlier studies. The DEF, CP, and 1YR are obtained from the Federal Reserve Statistical Release. Data on the dividend-price ratio, i.e., the dividend yield are available on Robert Shiller's website.[8] All data are monthly and cover the period April 1953 to November 2010.

In addition to the four financial variables, the analysis uses monthly and daily returns on the value-weighted market portfolio from the CRSP data files from April 1953 to December 2010. Monthly excess returns are calculated by subtracting the monthly yield on a three-month T-bill from the corresponding stock return. The three-month yield is used instead of the one-month yield because of the well-documented idiosyncrasies in this latter time series (see Duffee, 1996).

3.2. *Estimating Sharpe ratios*

The first step in the analysis is to determine if there is significant time-variation in monthly estimated conditional and realized Sharpe ratios, and, further, to see if the variation in these two series coincides. To compute the realized Sharpe ratio, we first calculate the realized volatility on a monthly basis using the sum of squared daily returns within the month:

$$v_t = \sqrt{\sum_{n=1}^{N} R_{n,t}^2}, \tag{6}$$

where v_t is the realized volatility for month t and $R_{n,t}$ are daily returns on the VW portfolio within the month. Adjusting for the daily, within-month mean return or subtracting the daily risk-free rate has no meaningful effect on the results. The realized volatility series is winsorized at the 99th percentile to eliminate a few obvious outliers. For example, the monthly realized volatilities in October 1987 and October 2008 exceed 23% (approximately 80% on an annualized basis). Moreover, these observations are more than eight standard deviations away from the mean of the series. The monthly realized Sharpe ratio for month t is then computed as

$$S_t \equiv \frac{R_t - R_{f,t-1}}{v_t}, \tag{7}$$

[8] http://www.econ.yale.edu/∼shiller/data.htm.

where R_t is the monthly return on the VW portfolio and $R_{f,t-1}$ is the corresponding monthly, risk-free rate at the beginning of the month.

We estimate the conditional Sharpe ratio using two different methodologies — (i) we estimate the first two conditional moments of returns separately and then take the ratio, and (ii) we estimate the conditional Sharpe ratio directly using a regression with the realized Sharpe ratio as the dependent variable. For the first methodology, expected returns are estimated by regressing excess returns on a vector of predetermined variables, and the conditional volatility is estimated by projecting realized volatility on to the same set of variables. Specifically, the conditional moments are modeled as follows:

$$E_t[R_{t+1} - R_{ft}] = X_t\beta_1, \tag{8}$$

$$SD_t[R_{t+1}] = X_t\beta_2, \tag{9}$$

where SD_t is the conditional standard deviation, and the corresponding regressions are specified as

$$R_{t+1} - R_{ft} = X_t\beta_1 + \varepsilon_{1t+1}, \tag{10}$$

$$v_{t+1} = X_t\beta_2 + \varepsilon_{2t+1}. \tag{11}$$

The conditioning variables are chosen based on the results in Whitelaw (1994), with the addition of lagged realized volatility as a conditioning variable in both equations. Specifically, we regress returns and realized volatility on the DEF, DIV, 1YR, CP, and lagged realized volatility (v_t).

Fitted values from Equations (10) and (11), can be used to compute conditional Sharpe ratios for each month. Specifically, based on information available at time t and parameter estimates from the estimation, the estimated conditional Sharpe ratio is

$$\hat{S}_{1,t} = \frac{E_t[R_{t+1} - R_{ft}]}{SD_t[R_{t+1}]} = \frac{X_t\hat{\beta}_1}{X_t\hat{\beta}_2}, \tag{12}$$

where the subscript 1 on the Sharpe ratio is to distinguish it from other estimates to be defined later.

Two alternative estimates of the conditional Sharpe ratio in Equation (12) shut down variation in one of the two conditional moments of returns, i.e.,

$$\hat{S}_{2,t} = \frac{\overline{R_{t+1} - R_{ft}}}{X_t\hat{\beta}_2} \tag{13}$$

$$\hat{S}_{3,t} = \frac{X_t \hat{\beta}_1}{\overline{v_{t+1}}}. \tag{14}$$

The conditional moment is replaced with the unconditional, in-sample mean of either excess returns or realized volatility. This estimate is equivalent to replacing the vector of conditioning variables in the regressions in Equations (10) and (11) with just a constant. In an out-of-sample context, these alternatives might outperform the more general specification due to overfitting or estimation error. In sample, they provide a decomposition of the estimated Sharpe ratio into its constituents.

Our second approach to estimating the conditional Sharpe ratio is to regress the realized Sharpe ratio from Equation (7) on the same set of predetermined variables used to estimate the conditional moments of returns above:

$$S_{t+1} = X_t \beta_3 + \varepsilon_{3t+1}. \tag{15}$$

The corresponding fitted conditional Sharpe ratio is

$$\hat{S}_{4,t} = X_t \hat{\beta}_3. \tag{16}$$

Thus, we now have four estimates in total: Equations (12)–(14) and (16).

3.3. *Empirical results*

Table 1, Panel A presents monthly results for the full sample period, May 1953 to December 2010 from the estimation of the conditional moments of returns based on the regressions in Equations (10) and (11). The results are broadly consistent with those reported in the literature. Both the dividend yield and the one-year Treasury rate are significant predictors of the market return at the 1% level. Lagged realized volatility has a negative, albeit small and statistically insignificant, coefficient. To the extent that this variable is a proxy for conditional volatility, this result coincides with the inability of many studies to find a significantly positive risk-return relation at the market level.[9] The *R*-squared of slightly less than 3% is lower than that reported in some previous studies, but the sample, which includes the recent financial crisis, may account for this result. (We examine this question in more detail below.)

In the volatility equation, lagged realized volatility, DEF, DIV, and CP are all highly significant. The significant positive serial correlation in realized

[9] See, for example, Guo and Whitelaw (2006) for an extended discussion and possible resolution of this puzzle.

Table 1. Estimation of the conditional first and second moments of returns.

	Const.	v_t	DEF	DIV	1YR	CP	R^2
			Panel A: Full Sample				
Mean	−0.194	−0.057	0.624	0.553***	−0.225***	−0.365	2.75%
	(0.687)	(0.140)	(0.642)	(0.181)	(0.081)	(0.728)	
Volatility	1.534***	0.585***	0.769***	−0.331***	−0.006	0.512***	54.37%
	(0.176)	(0.034)	(0.145)	(0.054)	(0.020)	(0.145)	
			Panel B: Pre-Crisis Sample				
Mean	−0.772	0.079	1.479**	0.505***	−0.340***	−0.056	4.01%
	(0.597)	(0.131)	(0.585)	(0.183)	(0.075)	(0.749)	
Volatility	1.600***	0.568***	0.412***	−0.309***	0.037*	0.471***	47.14%
	(0.165)	(0.039)	(0.143)	(0.051)	(0.021)	(0.148)	
			Panel C: Full Sample, Adjusted Dividend Yield				
Mean	1.623***	−0.150	0.577	0.781***	−0.199**	−0.031	2.97%
	(0.611)	(0.134)	(0.628)	(0.229)	(0.081)	(0.689)	
Volatility	0.677***	0.650***	0.603***	−0.180***	−0.030	0.260*	52.36%
	(0.126)	(0.036)	(0.148)	(0.059)	(0.020)	(0.150)	
			Panel D: Full Sample, Dividend Plus Repurchase Yield				
Mean	−0.833	−0.097	0.309	0.789***	−0.280***	−0.192	3.43%
	(0.712)	(0.135)	(0.659)	(0.190)	(0.081)	(0.695)	
Volatility	1.387***	0.632***	0.730***	−0.245***	−0.003	0.321**	52.99%
	(0.188)	(0.034)	(0.145)	(0.052)	(0.021)	(0.145)	

Note: Regressions of monthly, excess stock returns and realized volatilities for the CRSP VW index on lagged explanatory variables for the full sample, May 1953 to December 2010, and the pre-crisis sample, May 1953 to December 2007. The conditioning variables are lagged realized volatility (v_t); the Baa-Aaa spread (DEF); the dividend yield (DIV), the adjusted dividend yield, or the dividend plus repurchase yield; the one-year Treasury yield (1YR); and the commercial paper Treasury spread (CP). The model is given in Equations (10) and (11). Heteroscedasticity-consistent standard errors are in parentheses. Coefficients significant at the 10%, 5%, and 1% levels are marked with *, **, and ***, respectively.

volatility is a manifestation of the autoregressive conditional heteroskedasticity of monthly returns. Realized volatility is much more predictable than returns, with an R-squared exceeding 50%.

To examine the influence of the financial crisis on the regression results, we re-estimate the same specification ending the sample in December 2007, with the results reported in Table 1, Panel B. For this shorter sample, the coefficients are somewhat similar, but the default spread is now significant in the mean equation, and the one-year Treasury rate is significant at the 10% level in the volatility equation. Of perhaps greater interest, the R-squared in the mean equation increases to more than 4%. For the majority of the results that follow, we use the full sample, but it is important to keep in mind that explanatory power is reduced by the extreme and unpredictable variation of returns associated with the crisis and its aftermath.

There is some concern in the literature about structural instability of the mean return regression in Table 1 (e.g., Welch and Goyal, 2008). In particular, variables such as DIV appear to exhibit nonstationarity associated with structural shifts during the sample period, which may, in turn, account for instability in the regression coefficients estimated over shorter subsamples. Lettau and Van Nieuwerburgh (2008) investigate this issue and provide evidence of a shift in the mean of DIV to a lower level after 1991. To incorporate this evidence, we construct a new independent variable which is DIV adjusted for its mean within the two subsamples — up to an including 1991 and the period thereafter. Table 1, Panel C reports results for the regression over the full sample using this adjusted dividend yield (DIV-ADJ). The R-squared in the mean equation is somewhat higher than in the corresponding regression with the unadjusted dividend yield in Panel A, thus we chose this specification as our baseline, in-sample estimation.

The theoretical justification for the adjustment above would depend on the underlying mechanism behind the structural shift in the dividend yield. For example, if dividend yields declined because of a decline in risk aversion and a corresponding decline in the required risk premium (in a representative agent setting), then such an adjustment would be inappropriate. The declining dividend yield would reflect a decline in expected returns. However, in this case, the adjustment should not increase the explanatory power of the regression. Alternatively, the decline in the dividend yield could reflect a switch in corporate payout policy in favor of stock repurchases. Boudoukh *et al.* (2007) examine this issue in great detail and present evidence that adjusting for stock repurchases does increase the explanatory power of predictive regressions. As confirmation of this evidence Table 1, Panel D presents results for regressions with the dividend plus repurchase yield (DIV+REP) as the independent variable. Specifically, we use the annual DIV+REP series from Michael Roberts' website[10] to compute just the repurchase yield, i.e., the amount of stock repurchases divided by the market capitalization, at the market level. For each month in our sample, we take a weighted average of the relevant annual repurchase yields and add the number to our dividend yield series. For example, for the 12-month dividend yield that goes through March 1997, we add in 3/4 of the 1996 repurchase

[10] http://finance.wharton.upenn.edu/∼mrrobert/

yield and 1/4 of the 1997 repurchase yield.[11] In terms of R-squared, the results are a slight improvement over those in Panel C, which used the statistically adjusted DIV. Of greater importance, the essential implications of this analysis and the in-sample Sharpe ratio analysis that follows are insensitive to the precise specification.

Of course, all these results must be considered in the context of the well known problem with data snooping (see, for example Foster *et al.*, 1997). These concerns are mitigated somewhat by the fact that the predictor variables have been used in the literature for two decades, or more, and by the out-of-sample exercise conducted later in the paper.

Table 2 presents results of our second estimation methodology, i.e., direct estimation of the conditional Sharpe ratio, using DIV-ADJ described above for both the full sample and pre-crisis sample, although results for the original DIV series and DIV+REP are similar. For both samples, DIV and one-year Treasury rate are highly significant. The positive coefficient on the dividend yield is consistent with the fact that this variable positively predicts the return and negatively predicts volatility (see Table 1). The negative coefficient on the one-year interest rate reflects its negative relation to returns. The positive, albeit statistically insignificant, signs of the coefficients on DEF and CP are more difficult to reconcile with the earlier results given the positive relation between these variables and volatility. Finally, for the full sample, lagged realized volatility is a significant negative predictor of

Table 2. Direct estimation of the conditional Sharpe ratio.

	Const.	v_t	DEF	DIV-ADJ	1YR	CP	R^2
Full Sample	0.823*** (0.141)	−0.052** (0.024)	0.009 (0.124)	0.190*** (0.065)	−0.073*** (0.016)	0.100 (0.139)	3.71%
Pre-Crisis Sample	0.715*** (0.145)	−0.032 (0.027)	0.190 (0.141)	0.197*** (0.066)	−0.098*** (0.019)	0.140 (0.143)	4.35%

Note: Regressions of the realized Sharpe ratio for the CRSP VW index on lagged explanatory variables for the full sample, May 1953 to December 2010, and the pre-crisis sample, May 1953 to December 2007. The conditioning variables are lagged realized volatility (v_t), the Baa-Aaa spread (DEF), the adjusted dividend yield (DIV-ADJ), the one-year Treasury yield (lYR), and the commercial paper Treasury spread (CP). The model is given in Equation (15). Heteroscedasticity-consistent standard errors are in parentheses. Coefficients significant at the 10%, 5%, and 1% levels are marked with *, **, and ***, respectively.

[11] The annual DIV+REP series ends in 2003, so we assume a constant repurchase yield thereafter.

Table 3. Realized and conditional Sharpe ratios — descriptive statistics.

	Mean	Std Dev	AC(1)	Average Change	
				Peak-Trough	Trough-Peak
$\hat{\mu}_t$	0.549	0.836	0.876	1.193	−1.297
$\hat{\sigma}_t$	3.603	1.328	0.745	0.783	−0.590
$\hat{S}_{1,t}$	0.185	0.311	0.928	0.248	−0.332
$\hat{S}_{2,t}$	0.166	0.043	0.780	−0.026	0.013
$\hat{S}_{3,t}$	0.152	0.232	0.876	0.331	−0.360
$\hat{S}_{4,t}$	0.313	0.236	0.906	0.242	−0.277
S_{t+1}	0.311	1.180	0.037	1.703	−1.408

Correlations

	$\hat{\sigma}_t$	$\hat{S}_{1,t}$	$\hat{S}_{2,t}$	$\hat{S}_{3,t}$	$\hat{S}_{4,t}$	S_{t+1}
$\hat{\mu}_t$	−0.278	0.920	0.326	1.000	0.878	0.175
$\hat{\sigma}_t$		−0.358	−0.864	−0.278	−0.380	−0.076
$\hat{S}_{1,t}$			0.479	0.920	0.882	0.170
$\hat{S}_{2,t}$				0.326	0.460	0.072
$\hat{S}_{3,t}$					0.878	0.175
$\hat{S}_{4,t}$						0.200

Note: Descriptive statistics for the four, monthly, estimated conditional Sharpe ratio series ($\hat{S}_{1,t} - \hat{S}_{4,t}$), defined in Equations (12)–(14) and (16), respectively; the estimated conditional mean ($\hat{\mu}_t$) and volatility ($\hat{\sigma}_t$) of returns, defined in Equations (8) and (9); and the realized Sharpe ratio (S_{t+1}), defined in Equation (7). AC(1) is the first order autocorrelation. "Average Change" denotes the average difference between the series at successive peaks and troughs, or troughs and peaks, of the business cycle as determined by the NBER. All estimates use the adjusted dividend yield and are estimated over the full sample, May 1953 to December 2010.

the Sharpe ratio at the 5% level, which is consistent with the positive serial correlation in realized volatility documented above.

Table 3 provides descriptive statistics for the four estimated conditional Sharpe ratio series, for the realized Sharpe ratio, and for the first two conditional moments of returns. There are several notable results. First, variation in the estimated conditional mean appears to dominate the variation in the estimated conditional Sharpe ratios. The conditional mean is highly correlated with the conditional Sharpe ratio estimates in which both moments are allowed to vary, either explicitly or implicitly (#1, the ratio of the conditional moments, and #4, the direct estimation), but the magnitudes of the correlations between these series and the conditional volatility are much lower. Moreover, both the estimate using the ratio of the conditional moments (#1) and the directly estimated conditional Sharpe ratio

(#4) are highly correlated with the estimate in which the volatility is fixed at its sample average (#3), with correlations of approximately 0.9. In contrast, these same series have correlations of less than 0.5 with the estimate in which the mean is fixed at its sample average (#2). This result is not an artifact of the relative variation of the two conditional moments, since the conditional volatility has a substantially higher standard deviation. Of course, these are in-sample results. From an out-of-sample perspective, this phenomenon might be bad news if the in-sample explanatory power is due to overfitting of the conditional mean.

Second, all four estimated conditional Sharpe ratio series are positively correlated with the realized Sharpe ratio. For the three estimates that allow for variation in the conditional mean, these correlations are close to 0.2, somewhat more than twice the magnitude of the correlation with the estimate that fixes the conditional mean (#2). Perhaps surprisingly, the correlation for the estimate that fixes the conditional volatility (#3) is actually marginally higher than that for the estimate that allows both moments to vary over time (#1). These correlations are quite high given the sampling variation in the realized Sharpe ratio. This variation, due in large part to the unexpected components of returns and volatility, is apparent in the high standard deviation of the realized Sharpe ratio and its low first order autocorrelation.

Third, the means of the Sharpe ratio estimates calculated from the ratio of the individual moments (#1, #2, and #3) are approximately half as larger as the mean of either the realized Sharpe ratio or the direct estimate.[12] There is a Jensen's inequality effect, but the primary reason for the difference lies in the time series properties of returns and volatility. Specifically, the return and realized volatility series are relatively strongly negatively correlated, with a correlation of −0.27, a phenomenon that has been called the leverage effect.[13] These series themselves are made up of two components — an expected component that is reflected in the conditional moments series, and shocks. From Table 3, the correlation between the conditional moments is −0.28. This phenomenon is the well known and anomalous negative risk-return relation at the aggregate level.[14] It is this negative correlation that

[12] The means of these latter two series coincide by construction, i.e., the average fitted value equals the mean of the dependent variable in a regression context.

[13] If equity is a levered claim on underlying assets, a negative shock to the value of the assets, and therefore the equity, causes an increase in the volatility of the equity as it becomes more levered, for a fixed debt claim.

[14] See, for example, Whitelaw (1994) for a detailed discussion of this issue in a similar framework to that used in this paper.

causes the mean of the ratio of these moments (#1) to exceed the ratio of the means of the two moments, i.e., $0.18 > 0.55/3.60 = 0.15$. However, this effect is even stronger for the realized Sharpe ratio because there is substantially more variation in realized returns and volatility than in expected returns and conditional volatility. In addition, the shocks to the two moments are more negatively correlated than the moments, with a correlation between the residuals from Equations (10) and (11) of -0.33. This negative correlation between unexpected returns and shocks to volatility is consistent with a volatility feedback effect, which, in turn, is consistent with a positive risk-return tradeoff.[15]

A more formal way to address the relation between the estimated conditional Sharpe ratios and the realized Sharpe ratio is to run a regression of the latter on the former:

$$S_{t+1} = \alpha + \beta \hat{S}_{i,t} + \eta_{i,t+1}, \tag{17}$$

where i indexes the estimated conditional Sharpe ratio series. Running this regression for the direct estimation (#4) is meaningless, since the estimated conditional Sharpe ratio comes from a regression of the realized ratio on a set of predetermined variables. Therefore, such a regression will generate an intercept of zero, a slope coefficient of one, and an R-squared equal to that of the original regression. More generally, $\alpha = 0$ and $\beta = 1$ are the conditions for the forecast to be conditionally unbiased.

Table 4 presents estimation results for the regression in Equation (17) for the three series constructed as the ratio of the moments. The regressions are estimated over the full sample, including the financial crisis, and the forecasts are constructed using DIV-ADJ, although, as before, using an alternative dividend yield series generates qualitatively similar inferences. The key results are threefold. First, all the estimates have statistically significant forecasting power for the realized Sharpe ratio in that the hypothesis that the slope is zero can be rejected in all cases. Moreover, it is not possible to reject the hypothesis that the slope is equal to one in each case. Second, conditional unbiasedness can be rejected for series #1 and #3 due to the statistically significant intercept in the regression. This result is hardly surprising given the differences in means between the realized Sharpe ratio and the forecasts as documented in Table 3. Finally the R-squareds are relatively small, but, with the exception of the fixed mean estimate, they are

[15] See, for example, Guo and Whitelaw (2006) and Smith (2007) for illustrations of identifying the risk-return tradeoff using the volatility feedback effect.

Table 4. Regressions of the realized Sharpe ratio on conditional Sharpe ratios — in sample.

	α	β	R^2
$\hat{S}_{1,t}$	0.174***	0.740***	2.64%
	(0.055)	(0.180)	
$\hat{S}_{2,t}$	−0.014	1.966**	0.51%
	(0.164)	(0.991)	
$\hat{S}_{3,t}$	0.176***	0.894***	3.08%
	(0.052)	(0.178)	

Note: Regressions of the monthly realized Sharpe ratio for the CRSP VW index on three, monthly, estimated conditional Sharpe ratio series ($\hat{S}_{1,t} - \hat{S}_{3,t}$), defined in Equations (12)−(14), for the full sample, May 1953 to December 2010. All estimates use the adjusted dividend yield. Heteroscedasticity-consistent standard errors are in parentheses. Coefficients significant at the 10%, 5%, and 1% levels are marked with *, **, and ***, respectively.

comparable to the R-squared from the direct estimation (Table 2). In terms of in-sample explanatory power, fixing the mean has serious negative consequences, whereas fixing the variance has, if anything, a small positive effect on the forecasting.

Given the significant and predictable time-variation in stock market Sharpe ratios documented above, a natural question is how movements in estimated Sharpe ratios correspond to fluctuations in economic activity.[16] Figure 1 shows the estimated conditional Sharpe ratios and the NBER business cycle peaks and troughs (recessions, i.e., the peak-to-trough phases of the cycle, are marked by the shaded bars). There are only nine complete business cycles within the sample period; therefore, conclusions should be drawn with caution. Nevertheless, there appears to be striking cyclical variation in Sharpe ratios. Almost without exception, business cycle peaks correspond to low Sharpe ratios and business cycle troughs to high Sharpe ratios.

The last two columns in the top panel of Table 3 provide one quantification of this phenomenon. We calculate the difference between the conditional Sharpe ratio at the peak of the cycle and the subsequent trough, and then average these differences across the cycles in the sample. This average is a measure of how much the Sharpe ratio changes during the course of a recession. We also perform the same calculation from trough to peak, but, by

[16] Time-variations in both expected returns and volatility have been previously linked to the business cycle. See, for example, Fama and French (1989) and Schwert (1989).

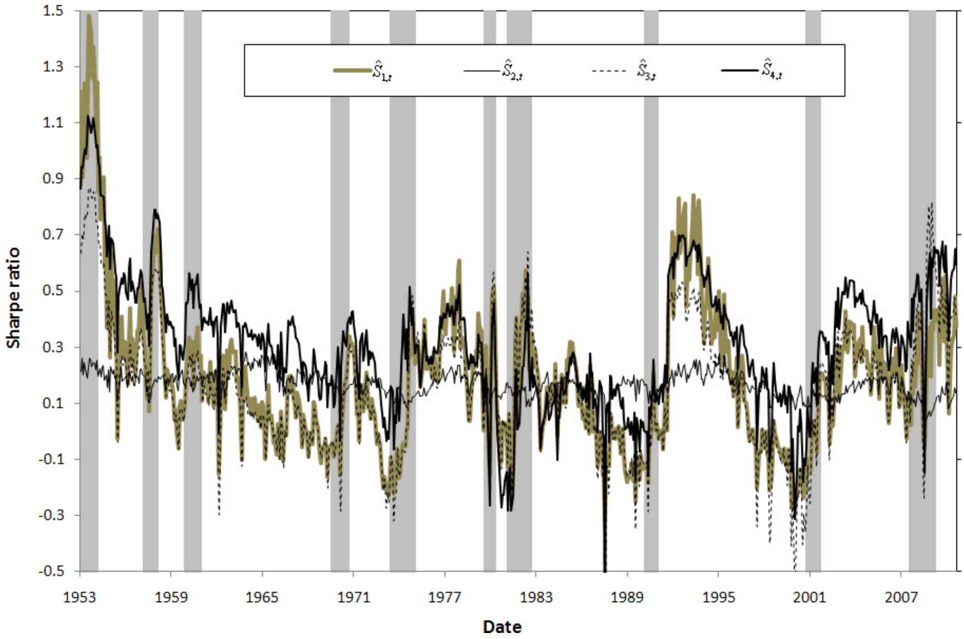

Fig. 1. Estimated conditional Sharpe ratios.
Note: Monthly expected conditional Sharpe ratios ($\hat{S}_{1,t} - \hat{S}_{4,t}$), defined in Equations (12)–(14) and (16), respectively. All estimates use the adjusted dividend yield and are estimated over the full sample, May 1953–December 2010. NBER recessions are marked by shaded bars.

construction, this average difference is approximately the negative of the change from peak to trough. The time from peak to trough (i.e., the contractionary phase of the cycle) is short, averaging less than 12 months, but the average increase in the monthly estimated rolling Sharpe ratio is large, on the order of 0.25 for our unconstrained estimators (#1 and #4), which would be greater than 0.8 on an annualized basis. Interestingly, both the conditional mean and volatility increase from peak to trough, but the former effect dominates.

The data indicate that the return/volatility tradeoff is significantly better entering expansions than it is leaving expansions. To some extent, this result is consistent with the theoretical results in Whitelaw (2000). At the end of expansions, when the probability of shifting to a contraction is high, the conditional volatility is also high. However, equity returns depend on whether a regime switch occurs, an event that is independent of the marginal rate of substitution. Consequently, the correlation in Equation (5) is low and so is the expected return. Clearly, this combination will yield the low Sharpe ratios shown in Figure 1. The regime-shift model also generates a similar prediction for transitions from contractions to expansions, but a decrease in Sharpe ratios at this point in the cycle is not evident in the data. The one mitigating factor in the model in Whitelaw (2000) is that regime switch probabilities are much more stable within contractions. Consequently, the volatility and expected return effects are smaller.

Figure 1 has several other notable features. With the exception of the fixed-mean estimate, the conditional Sharpe ratios appear to coincide, consistent with their high correlations in Table 3. The fixed-mean estimate shows substantially less variation, consistent with its low standard deviation in Table 3, and there is no strong business cycle pattern. The conditional Sharpe ratios also go negative at various points in the sample, again with the exception of the fixed mean series, i.e., the expected return on the market is less than the risk-free rate. This result is somewhat puzzling, although negative risk premiums are not theoretically precluded in the framework of Section 2 (see Boudoukh *et al.*, 1997). Moreover, these results are consistent with the results in Kairys (1993), who uses commercial paper rates to predict negative risk premiums.

4. Exploiting Predictable Variation

We now turn to an out-of-sample analysis of predictable variation in stock market Sharpe ratios, using both rolling and expanding window regressions. After examining the properties of the estimated conditional Sharpe ratios and evaluating their forecasting power, we employ these forecasts to construct simple market-timing strategies.

4.1. *Out-of-sample forecasting*

The previous section documents statistically significant time-variation in conditional Sharpe ratios, and statistically and economically significant predictive power for estimates based on a simple linear model. From a practical perspective, however, the key issue is whether the empirical model has economically significant out-of-sample predictive power. Unfortunately, it is difficult to conduct a true out-of-sample test. The conditioning variables are chosen based on their correlation with returns and volatility in a sample that runs through April 1989, leaving over 20 years of new data, but there is still the issue of how choices about which papers to write and publish create their own data snooping problems. Nevertheless, it is worthwhile to consider the predictive power of out-of-sample regressions.

There are two possible ways to specify the out-of-sample regressions for estimating the conditional Sharpe ratios. The first method is to choose a fixed sample size and to run rolling regressions. That is, a fixed number of observations are used to estimate each set of coefficients, and the estimation window is moved forward by one month at a time. The advantage of this approach is that if the coefficients vary over time, either because of misspecification of the empirical model or structural shifts, then the coefficients from the rolling regressions will "adapt" to these changes. The second method is to add the new monthly observation to the estimation dataset as we move through time. As a result, the number of observations increases through time, and the later coefficients of these cumulative regressions will be less subject to estimation error if the empirical specification is correct.

An additional issue is the choice of DIV series. The DIV-ADJ series used in the in-sample analysis relies on an in-sample estimation of the break point in the original series (Lettau and Van Nieuwerburgh, 2008). Instead, for the out-of-sample analysis, we use the DIV+REP series.

The most natural way to evaluate these alternatives is to examine their out-of-sample performance. For both the rolling and cumulative regressions, the initial estimation period is chosen to be 10 years, i.e., 120 monthly observations from May 1953 to April 1963 are used to estimate the first set of coefficients. These coefficients and the data on the explanatory variables from April 1963 are then used to estimate the conditional moments of stock market returns for May 1963. The estimation is then rolled forward one month, adding the most recent observation for both the rolling and cumulative regressions, but also dropping the oldest observation from the rolling regression. Both techniques generate a series of 572 out-of-sample conditional Sharpe ratios.

Before turning to the out-of-sample forecasting power, one potentially important issue that can be addressed using the rolling regressions is the instability of the coefficient estimates. Unstable coefficients indicate either structural shifts in the data, a misspecified model, or significant estimation error. In the case of structural shifts, predictive power might be gained from shortening the estimation period further, although there is a clear tradeoff with estimation error as the number of observations decreases. In the case of model misspecification, alternative specifications, specifically more flexible functional forms, might prove useful, but these formulations will also likely increase estimation error. Figure 2 shows the rolling coefficient estimates for the conditional mean (Panel A), the conditional volatility (Panel B), and the direct estimation of the Sharpe ratio (Panel C). Note that the date on the x-axis refers to the last date in the estimation period, e.g., the coefficient for

December 1978 is based on the ten years of data from January 1969 to December 1978.

The graphs show a good deal of instability. In the mean equation, consider the two variables that are significant in the in-sample regression — DIV+ REP and the one-year Treasury rate (Table 1, Panel D). The coefficient on DIV+REP ranges up to seven, with a small period of negative coefficients associated with rise of the technology bubble in the mid to late 1990s. The coefficient on the interest rate is predominantly negative, but there are periods in the late 1990s and during the financial crisis when it is positive. From an economic standpoint, the switching of the sign of the coefficient is especially worrisome, but one must keep in mind the substantial sampling error associated with a 10-year estimation window.

In the volatility equation, all the variables but the Treasury rate are significant in the in-sample regression (Table 1, Panel D). The coefficient on

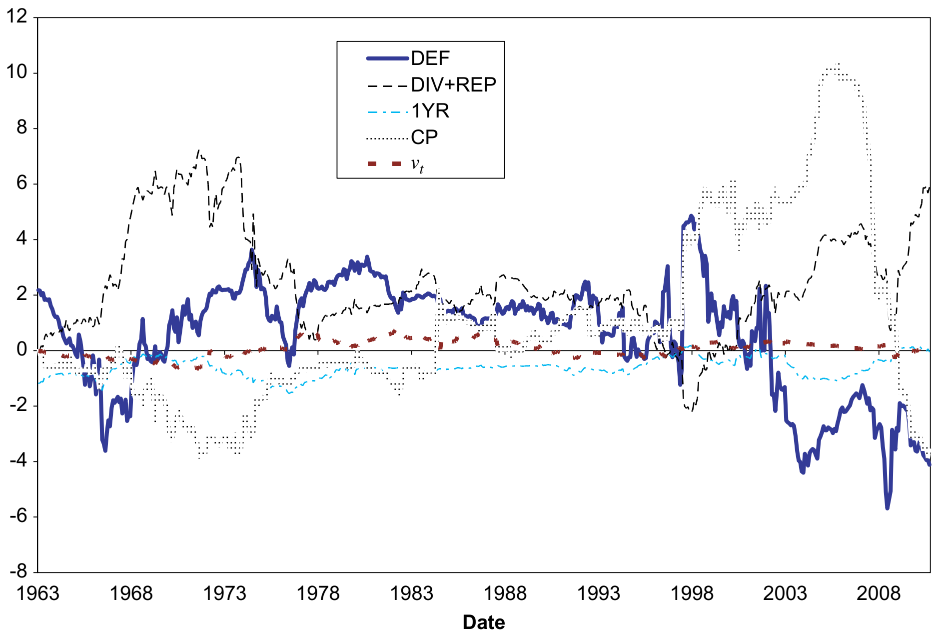

(a) Panel A: Conditional Mean

Fig. 2. Estimated rolling regression coefficients.
Note: Estimated coefficients from l0-year rolling regressions for the conditional mean (Panel A), the conditional volatility (Panel B) Equations (10) and (11), and the direct estimation of the conditional Sharpe ratio (Panel C) Equation (15) of the CRSP VW portfolio. All estimates use the dividend plus repurchase yield and cover the periods from May 1953−April 1963 through January 2001−December 2010.

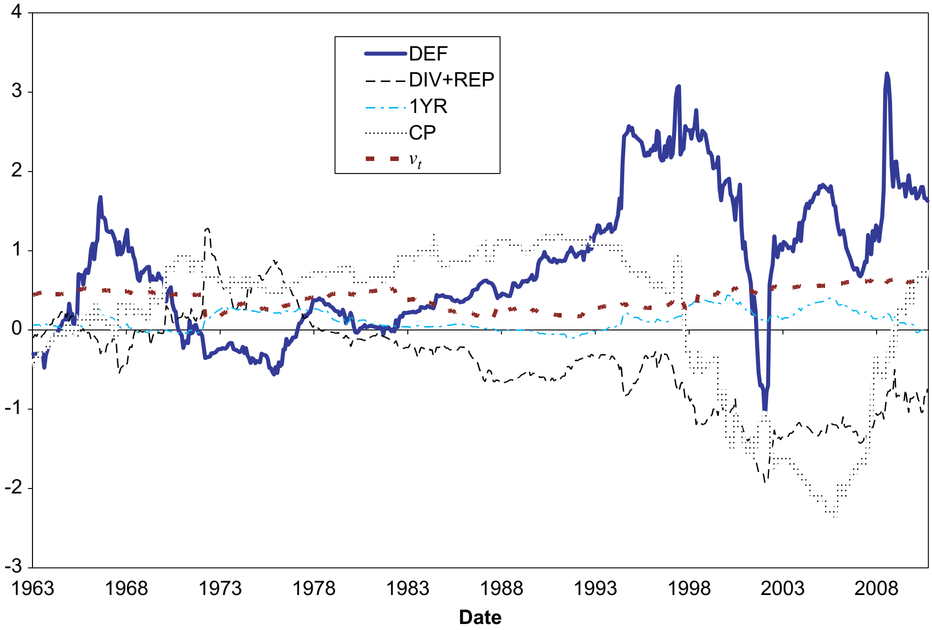

(b) Panel B: Conditional Volatility

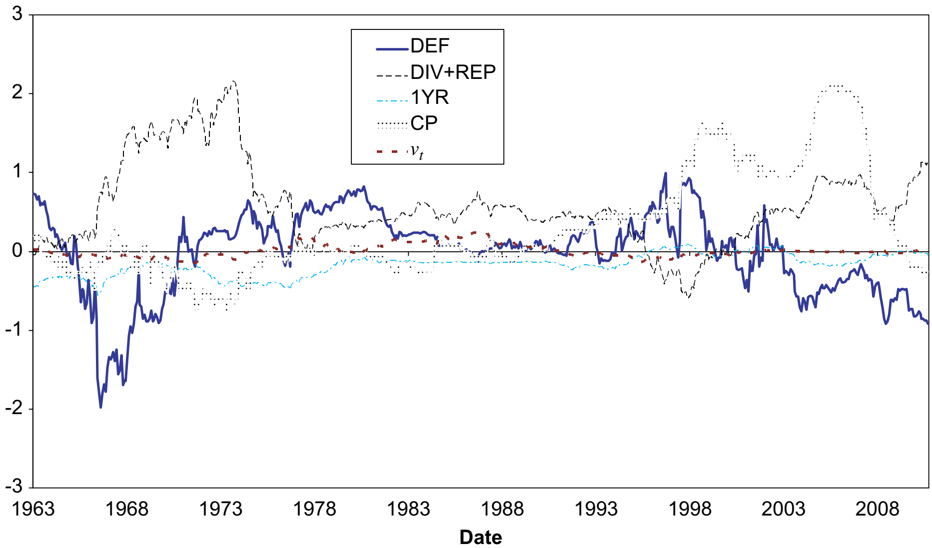

(c) Panel C: Conditional Sharpe Ratio (Direct Estimation)

Fig. 2. (*Continued*)

lagged realized volatility is clearly the most stable, illustrating the robustness of the persistence of volatility. In contrast, the coefficients on the default spread, DIV+REP, and CP switch signs at least twice, at different points in the sample. Again, however, these results should be considered in the context of a regression over a short sample period with substantial estimation error.

Direct estimation of the Sharpe ratio does not appear to improve the stability of the rolling 10-year coefficients. The two significant variables from the in-sample regression in Table 2, DIV+REP and the Treasury rate, show patterns that are remarkably similar to those from the estimation of the conditional mean. Given these variations, any out-of-sample forecasting power of the rolling regression might perhaps be even more surprising.

Table 5 addresses the issue of out-of-sample forecasting power directly. Panel A presents results from the regression of the realized Sharpe ratio on the estimated conditional Sharpe ratios for the period May 1963 to December 2010, using both the rolling window and expanding window regressions described above. These regressions are comparable to the in-sample regressions in Table 4, except that they omit the first 10 years of the sample due to the necessity of using this window to estimate the first set of coefficients.

There is definite evidence of forecasting power. All the slope coefficients are positive, i.e., the realized Sharpe ratio is positively related to the estimated conditional Sharpe ratio, and three of the four are statistically significant for the rolling regressions. While the magnitudes of the slope coefficients are comparable for the expanding window regressions, only two of the four coefficients are significant at the 10% level. The R-squareds for the rolling regressions reach just over 1%, which is only slightly greater than one-third of the magnitude in the comparable in-sample regressions.

Of the four estimated Sharpe ratios, statistical evidence of forecasting power is weakest for the estimate that fixes the mean return at its value in the estimation window. While the slope coefficient is largest for this series, the magnitude reflects the lack of variation in the series. For neither the rolling nor expanding window regressions is the coefficient statistically significant, and the R-squareds are less than 0.5%. This result coincides with the in-sample evidence, i.e., time-variation in the conditional expected return appears to be the dominant component of variation in the Sharpe ratio.

Finally, there is also substantial evidence of overfitting across all the series. The hypothesis that the forecast is conditionally unbiased, i.e., an intercept of zero and a slope of one, can be rejected for a significant majority of the series, and the failure to reject for the others is due to lack of power, i.e., large

Table 5. Out-of-sample analysis.

	Rolling			Expanding		
	α	β	R^2	α	β	R^2
Panel A: Regressions of the Realized Sharpe Ratios on Estimated Conditional Sharpe Ratio						
$\hat{S}_{1,t}$	0.219***	0.257**	0.99%	0.226***	0.238	0.50%
	(0.050)	(0.106)		(0.052)	(0.148)	
$\hat{S}_{2,t}$	0.155**	0.638	0.42%	0.173	0.405	0.12%
	(0.078)	(0.419)		(0.107)	(0.521)	
$\hat{S}_{3,t}$	0.224***	0.237***	1.21%	0.224***	0.230**	0.71%
	(0.049)	(0.084)		(0.051)	(0.116)	
$\hat{S}_{4,t}$	0.224***	0.063***	1.07%	0.223***	0.072*	0.62%
	(0.049)	(0.023)		(0.052)	(0.039)	

	Rolling	Expanding
Panel B: Predictive R-Squareds		
$\hat{S}_{1,t}$	−8.48%	−6.05%
$\hat{S}_{2,t}$	−0.47%	−0.37%
$\hat{S}_{3,t}$	−12.66%	−8.48%
$\hat{S}_{4,t}$	−233.90%	−105.25%

Note: Panel A presents results from regressions of the monthly realized Sharpe ratio for the CRSP VW index on four, monthly, estimated conditional Sharpe ratio series ($\hat{S}_{1,t} - \hat{S}_{4,t}$), defined in Equations (12)−(14) and (16), estimated over rolling, 10-year windows. All estimates use the dividend plus repurchase yield and cover the periods from May 1953−April 1963 through January 2001−December 2010. Heteroscedasticity-consistent standard errors are in parentheses. Coefficients significant at the 10%, 5%, and 1% levels are marked with *, **, and ***, respectively. Panel B presents predictive R-squareds, as defined in Equation (18), for the same four series.

standard errors. Moreover, all the slope coefficients are less than one. This overfitting is illustrated clearly in Table 5, Panel B. This panel reports what we call the predictive R-squared for each forecast series. Specifically, we calculate an R-squared measure imposing conditional unbiasedness:

$$R^2 = 1 - \frac{\sum (S_{t+1} - \hat{S}_{i,t})^2}{\sum (S_{t+1} - \mu_S)^2}, \tag{18}$$

where μ_S is the mean realized Sharpe ratio. The predictive R-squareds in Panel B are uniformly negative. In other words, while the conditional Sharpe ratios have forecasting power for the realized Sharpe ratio, the unconditional mean provides a better fit than does the conditional Sharpe ratio.

4.2. *Trading strategies*

While the above results clearly demonstrate economically and statistically significant out-of-sample forecasting power for the realized Sharpe ratio, they are not linked directly to a feasible trading strategy. The possible set of trading strategies is huge; nevertheless, it is worthwhile to look at the performance of a few stylized strategies. Consider the strategy of estimating the conditional Sharpe ratio using the 10-year rolling regression and comparing this number to a fixed threshold. If the estimated conditional Sharpe ratio is larger than the threshold, then invest in the stock market; if it is smaller, then invest in the risk-free asset. We then consider the Sharpe ratio of the months when the strategy is invested in equities. This ratio can easily be compared to a buy-and-hold strategy that always holds the market.[17]

Table 6 reports the results from executing four strategies: a buy-and-hold strategy and three market-timing strategies where the thresholds for investing in the market are three different pre-specified conditional Sharpe ratio levels — 0.0, 0.1, and 0.2. Again, we consider all four of our conditional Sharpe ratio series and we use DIV+REP throughout. The third column of the table gives the number of months, out of a possible 572, in which the strategy is invested in the market. The table also shows the mean and average realized volatility of monthly stock market returns for the months in which the market is held. For the buy-and-hold strategy, these are the sample averages for the full time period. The last two columns present statistics calculated from monthly returns (rather than the daily returns that are used to compute realized volatility). The ex post Sharpe ratio is the mean excess monthly return over the volatility of the return when the strategy is in the market, while ex post volatility is just the denominator of this ratio.

[17]These market-timing strategies ignore both transaction costs and information in the magnitude of the conditional Sharpe ratios relative to the threshold. A more sophisticated, and potentially better performing, strategy might involve time-varying market weights that depend on the prior position in the market and the relative magnitude of the conditional Sharpe ratio. Nevertheless, the stylized strategy is sufficient to illustrate the extent of predictable variation.

Table 6. Trading strategies.

	Signal	No. of Months	Mean Return	Avg. Realized Vol.	Avg. Realized SR	Ex post SR	Ex post Vol.
Full sample		572	0.448	3.786	0.252	0.099	4.522
$\hat{S}_{1,t}$	CSR > 0.0	326	0.827	3.870	0.328	0.184	4.507
	CSR > 0.1	267	0.955	3.920	0.350	0.209	4.562
	CSR > 0.2	207	1.111	3.967	0.365	0.250	4.441
$\hat{S}_{2,t}$	CSR > 0.0	493	0.332	3.627	0.248	0.078	4.268
	CSR > 0.1	384	0.422	3.357	0.285	0.104	4.045
	CSR > 0.2	181	0.400	2.756	0.351	0.113	3.539
$\hat{S}_{3,t}$	CSR > 0.0	326	0.827	3.870	0.328	0.184	4.507
	CSR > 0.1	269	0.961	3.930	0.353	0.207	4.639
	CSR > 0.2	210	1.082	4.178	0.367	0.225	4.811
$\hat{S}_{4,t}$	CSR > 0.0	326	0.827	3.870	0.328	0.184	4.507
	CSR > 0.1	301	0.873	3.888	0.337	0.192	4.550
	CSR > 0.2	289	0.872	3.925	0.335	0.189	4.609

Note: Descriptive statistics for trading strategies that hold the market portfolio when the estimated conditional Sharpe ratio exceeds the given threshold. The estimated conditional Sharpe ratio series ($\hat{S}_{1,t} - \hat{S}_{4,t}$), are defined in Equations (12)−(14) and (16) and estimated over rolling, 10-year windows. All estimates use the dividend plus repurchase yield and cover the periods from May 1953−April 1963 through January 2001−December 2010. The mean return and average realized volatility and Sharpe ratio are for those months when the strategy is long the market. The ex post Sharpe ratio is the Sharpe ratio calculated using monthly returns for these months, and the ex post volatility is the volatility of these monthly returns.

There are several notable results. First, there is consistent evidence of forecasting power. Of the 12 strategies (four predictors × three thresholds), eleven have Sharpe ratios that exceed that of the buy-and-hold. The one exception is the fixed mean predictor with a threshold of 0. Moreover, across the four signals, three produce monotonically increasing Sharpe ratios as the threshold increases. For the fourth, the direct estimation of the Sharpe ratio, there is only a very small violation of monotonicity. Perhaps of greatest importance, the increase in Sharpe ratio associated with the strategies is economically large. For example, for the estimate constructed as the ratio of the conditional moments, where both moments vary over time, the Sharpe ratios of the strategies exceed that of the buy-and-hold by 30−45%. These strategies are invested in the stock market in 36−57% of the months in the sample.

Second, for the strategies that allow the expected return to vary over time (#1 and #3), the forecasting power is coming primarily from this variation. Mean returns increase monotonically with the threshold, and the average realized volatility varies little, increasing as the Sharpe ratio increases if anything. Interestingly, the fixed mean strategy is able to identify periods with low realized volatilities and hence high realized Sharpe ratios. However, this predictive power is much less marked for volatilities and Sharpe ratios calculated using monthly returns. For the direct estimation of the Sharpe ratio, the threshold appears to be the least important. The number of months invested changes less as the threshold varies and the performance statistics also vary little.

Overall, the results confirm the earlier conclusions from both the in-sample and out-of-sample analyses, i.e., that there is economically significant, predictable variation in stock market Sharpe ratios.

5. Conclusion

This paper demonstrates the ability of relatively straightforward linear specifications, of either the conditional mean and volatility of equity returns or of the Sharpe ratio directly, to predict dramatic time-variation in monthly, stock market Sharpe ratios. This predictability is evident both in-sample and out-of-sample, where market-timing strategies outperform a buy-and-hold strategy in terms of ex post Sharpe ratios. This evidence provides further support for the contention that the mean and volatility of stock market returns do not move together.

Variations of the magnitude documented are inconsistent with the conditional CAPM and related models that imply a close to constant market Sharpe ratio. One possible explanation is that the results are due to market irrationality or inefficiency. However, the apparent relation between variation in Sharpe ratios and the business cycle suggests the possibility of an economic interpretation. Whitelaw (2000) provides a rational expectations, general equilibrium model that is broadly consistent with the empirical evidence. In this model, discrete shifts between expansions and contractions overturn the standard return/volatility relation. Alternatively, large fluctuations in risk aversion, as in Campbell and Cochrane (1999), could also account for significant time-variations in Sharpe ratios. Further research in this area is warranted.

While the empirical evidence provides insights into the time series properties of equity returns and their underlying economics, it also has implications in other areas. In particular, substantial, predictable time-variation in market Sharpe ratios casts doubt on the ability of the volatility of even broadly diversified portfolios to proxy for priced risk. Consequently, standard measures of investment performance and traditional portfolio asset allocation rules may have to be re-thought.

Acknowledgments

Thanks to Kobi Boudoukh, Michael Brennan, Ravi Jagannathan, and Matt Richardson for helpful comments.

References

Bansal, R., and S. Viswanathan, 1993, No Arbitrage and Arbitrage Pricing, *Journal of Finance* 48, 1231−1262.

Boudoukh, J., R. Michaely, M. Richardson and M. Roberts, 2007, On the Importance of Measuring Payout Yield: Implications for Empirical Asset Pricing, *Journal of Finance* 62, 877−915.

Boudoukh, J., M. Richardson and R. Whitelaw, 1997, Nonlinearities in the Relation Between the Equity Risk Premium and the Term Structure, *Management Science* 43, 371−385.

Breeden, D., M. Gibbons and R. Litzenberger, 1989, Empirical Tests of the Consumption-Oriented CAPM, *Journal of Finance* 44, 221−262.

Breen, W., L. Glosten and R. Jagannathan, 1989, Economic Significance of Predictable Variations in Stock Index Returns, *Journal of Finance* 44, 1177−1189.

Campbell, J., 1987, Stock Returns and the Term Structure, *Journal of Financial Economics* 18, 373−399.

Campbell, J., and J. Cochrane, 1999, By Force of Habit: A Consumption-Based Explanation of Aggregate Stock Market Behavior, *Journal of Political Economy* 107, 205−251.

Chen, N.-F., R. Roll and S. Ross, 1986, Economic Forces and the Stock Market, *Journal of Business* 59, 383−404.

Duffee, G., 1996, Idiosyncratic Variation of Treasury Bill Yields, *Journal of Finance* 51, 527−551.

Fama, E., and K. French, 1989, Business Conditions and Expected Returns on Stocks and Bonds, *Journal of Financial Economics* 25, 23−49.

Fama, E., and K. French, 1993, Common Risk Factors in the Returns on Stocks and Bonds, *Journal of Financial Economics* 33, 3−56.

Ferson, W., and C. Harvey, 1992, Seasonality and Consumption-Based Asset Pricing, *Journal of Finance* 47, 511−552.

Fleming, J., C. Kirby and B. Ostdiek, 2001, The Economic Value of Volatility Timing, *Journal of Finance* 56, 329−352.

Foster, D., T. Smith and R. Whaley, 1997, Assessing Goodness-of-Fit of Asset Pricing Models: The Distribution of the Maximal R^2, *Journal of Finance* 52, 591−607.

French, K., W. Schwert and R. Stambaugh, 1987, Expected Stock Returns and Volatility, *Journal of Financial Economics* 19, 3−29.

Glosten, L., R. Jagannathan and D. Runkle, 1993, On the Relation Between the Expected Value and the Volatility of the Nominal Excess Return on Stocks, *Journal of Finance* 48, 1779−1801.

Guo, H., and R. Whitelaw, 2006, Uncovering the Risk-Return Relation in the Stock Market, *Journal of Finance* 61, 1433−1463.

Hansen, L., and K. Singleton, 1983, Stochastic Consumption, Risk Aversion, and the Temporal Behavior of Asset Returns, *Journal of Political Economy* 91, 249−265.

Harrison, M., and D. Kreps, 1979, Martingales and Arbitrage in Multiperiod Security Markets, *Journal of Economic Theory* 20, 381−408.

Harvey, C., and A. Siddique, 2000, Conditional Skewness in Asset Pricing Tests, *Journal of Finance* 55, 1263−1295.

Kairys, J., 1993, Predicting Sign Changes in the Equity Risk Premium Using Commercial Paper Rates, *Journal of Portfolio Management* 20, 41−51.

Kandel, S., and R. Stambaugh, 1990, Expectations and Volatility of Consumption and Asset Returns, *Review of Financial Studies* 3, 207−232.

Kandel, S., and R. Stambaugh, 1996, On the Predictability of Stock Returns: An Asset Allocation Perspective, *Journal of Finance* 51, 385−424.

Keim, D., and R. Stambaugh, 1986, Predicting Returns in the Stock and Bond Markets, *Journal of Financial Economics* 17, 357−390.

Kraus, A., and R. Litzenberger, 1976, Skewness Preference and the Valuation of Risky Assets, *Journal of Finance* 31, 1085−1100.

Lettau, M., and S. Van Nieuwerburgh, 2008, Reconciling the Return Predictability Evidence, *Review of Financial Studies* 21, 1607−1652.

Lucas, R., 1978, Asset Prices in an Exchange Economy, *Econometrica* 46, 1429−1445.

Merton, R., 1973, An Intertemporal Capital Asset Pricing Model, *Econometrica* 41, 867−887.

Schwert, W., 1989, Why Does Stock Market Volatility Change Over Time? *Journal of Finance* 44, 1115−1153.

Smith, D., 2007, Risk and Return in Stochastic Volatility Models: Volatility Feedback Matters! Unpublished working paper, Queensland University of Technology.

Welch, I., and A. Goyal, 2008, A Comprehensive Look at the Empirical Performance of Equity Premium Prediction, *Review of Financial Studies* 21, 1455−1508.

Whitelaw, R., 1994, Time Variations and Covariations in the Expectation and Volatility of Stock Market Returns, *Journal of Finance* 49, 515−541.

Whitelaw, R., 2000, Stock Market Risk and Return: An Equilibrium Approach, *Review of Financial Studies* 13, 521−547.