Bocconi

WORKING PAPER SERIES

WORKING PAPER
N. 37
DECEMBER 2016

DO REGIMES IN EXCESS STOCK RETURN PREDICTABILITY CREATE ECONOMIC VALUE? AN OUT-OF-SAMPLE PORTFOLIO ANALYSIS

Giulia Dal Pra, Massimo Guidolin, Manuela Pedio and Fabiola Vasile

Do Regimes in Excess Stock Return Predictability Create Economic Value? An Out-of-Sample Portfolio Analysis

Giulia Dal Pra
Bocconi University

Massimo Guidolin†
Bocconi University and Baffi CAREFIN

Manuela Pedio
Bocconi University and Baffi CAREFIN

Fabiola Vasile
Euromobiliare SGR

Abstract

We analyze the recursive, out-of-sample performance of asset allocation decisions based on financial ratio-predictability under single-state linear and regime-switching models. We adopt both a statistical perspective to analyze whether models based on the dividend-price, earning-price, and book-to-market ratios can forecast excess equity returns, and an economic approach that turns predictions into portfolio strategies. The strategies consist of a portfolio switching approach, a mean-variance framework, and a long-run dynamic model. We report an interesting disconnect between a statistical perspective, whereby the ratios yield a modest forecasting power, and a portfolio approach, by which a moderate predictability is occasionally sufficient to yield significant portfolio outperformance, especially before transaction costs and when regimes are taken into account. However, also when regimes are considered, predictability gives high payoffs only to long-horizon, highly risk-averse asset managers. Moreover, different strategies deliver different performance rankings across predictors. Finally, we find evidence inconsistent with the notion that increasing sophistication in the way portfolio decisions are modeled, delivers a superior performance.

Key words: predictability, Markov switching, economic value, optimal portfolio choice.

1. Introduction

Whether excess equity returns may be predictable using simple (aggregate) financial ratios represents one of the classical debates in the literature and appears to be rooted in the early

† Corresponding author. Address: Bocconi University, Finance dept., via Roetgen 1, 20136 Milan, Italy. E-mail: massimo.guidolin@unibocconi.it. Tel.: +39-02-5836.3556.

attempts to test the celebrated (weak-form) efficient markets hypothesis, i.e., the claim that trading prices would contain all relevant information, thus making impossible to achieve any risk-adjusted profit exploiting only public information, see e.g., the review in Rapach and Zhou (2012). As of today, the literature has still not reached an agreement, with empirical papers offering conclusions that range from stark rejections of the hypothesis of *statistically* predictable excess returns on the basis of simple financial ratios (see, e.g., Bossaerts and Hillion, 1999; Welch and Goyal, 2008) to a less sanguine view that leaves the possibility open that for some data sets and predictors, and under appropriate constraints and restrictions on the resulting allocations, predictability may exist and be exploitable (see e.g., Campbell and Thompson, 2008, Cardinale et al., 2014; Cochrane, 2008; Famy, 2007; Guidolin et al., 2009).

This debate has often emphasized the need to take steps beyond a discussion of whether excess asset returns would be statistically predictable, to tackle the issue of the economic value associated to such a weak predictability (see, e.g., Bulla et al., 2011; Campbell and Thompson, 2008; Cardinale et al., 2014; Fugazza et al., 2009; Pesaran and Timmermann, 1995). Moreover, a strand of the literature exists that has shown that capturing the instability in standard predictive regressions—for instance, through models that are included in the standard toolkit available to portfolio managers, such as regime switching models (see e.g., Ang and Bekaert, 2004; Ang and Timmermann, 2011; Bulla et al., 2011; Fabozzi et al., 2006; Guidolin and Ria, 2011; Henkel et al., 2011; Kritzman et al., 2012)—may not only increase the (in- and out-of-sample) predictability of excess stock returns, but also generate non-negligible economic value, in the form of superior risk-adjusted realized performances. In this paper, we contribute to this debate by systematically examining whether, when, and how a few standard valuation ratios may predict excess equity returns and generate risk-adjusted outperformance when we take into account that most predictive relationships are subject to statistical instability and appear to be elusive even to relatively experienced analysts. In addition, we investigate of a few of the boundaries—here represented by the complexity of the asset allocation models, the inclusion of transaction costs, and the length of the investment horizon when periodic rebalancing is possible—that may delimit the "space" of effectiveness of regimes in portfolio choice.

The interest of the portfolio management literature for statistical models featuring alternative regimes is relatively unsurprising. Financial markets change their behavior over time following bull and bear cycles that shape the distribution of excess asset returns. Contrary to standard jumps, which are transitory and sudden price movements, these shifts can be characterized by such a persistence to the point of making it compelling to consider them in asset allocation decisions.[2] In particular, standard predictive regressions may be unable to fit the properties of most return series, because these are often generated by more complicated models that reflect the presence of persistent shifts in the strength with which predictors affect excess returns (see Ang and Timmermann, 2010, and Guidolin, 2011b). The Markov switching models used in this paper, allow us to perform inference on structural breaks and to capture the dynamics of financial returns.

In particular, in this paper we investigate under which conditions it is possible to forecast excess stock returns through three alternative aggregate financial indicators—the dividend-price (henceforth, DP), the earning-price (EP), and the book-to-market (BM) ratios—in a way that allows investors to obtain valuable risk-adjusted performances.[3] To perform this exercise, we resort to three of the most robust predictors that the previous literature has isolated (see Rapach and Zhou, 2013, and Welch and Goyal, 2008), with reference to a standard set of US monthly data for the sample 1926-2012. The novelty of our work consists in recognizing that the very strength of the statistical predictability may be time-varying according to a regime-switching process driven by a standard first-order Markov chain. Our analysis is then articulated on the contrast between the weakness of the statistical predictability characterizing the data in the case of simple regression models and the likelihood that sizeable risk-adjusted profits may be generated by adopting a portfolio strategy that exploits not only the signals provided by the financial ratios, but also those coming from an estimated hidden Markov state.

The breadth of the effort that we perform represents one of the key features of our work. Our results derive from a range of statistical models that include simple regressions that use predictors on the right-hand side, Markov switching (MS) regressions that imply time-varying slope coefficients associated to the predictors, and Markov Switching Vector Autoregressions (MSVAR) that endogenize the dynamics of the predictor variables. Moreover, we perform an analysis of three different asset allocation frameworks. First, we implement simple switching strategies à la Pesaran

---

[2] See Das and Uppal (2003) or Guidolin and Ossola (2009) on whether accounting for standard jumps in optimal asset allocation has a chance to cause first-order impacts.

[3] In this paper we only use one predictor at the time. Indeed, as discussed by Welch and Goyal (2008) and Rapach and Zhou (2013) this approach offers identical or even stronger predictive performances as models that employ all predictors at the same time, in what is called a "kitchen-sink" design. This is presumably due to considerable problems of multicollinearity and parameter proliferation that tend to plague larger models.

and Timmermann (1995) that best characterize risk-neutral investors, in which an asset manager goes long in stocks only when the predicted risk premium is positive. Second, we adopt static mean-variance (henceforth, MV) strategies that characterize risk-averse investors, that assumes that higher variances ought to be compensated by higher means, but ignore higher moments of excess return distribution. Finally, we discuss dynamic strategies in which a risk-averse investor with power utility preferences accounts for the effects of predictability and regimes on the forecast density of excess returns. In the case of switching strategies, when it is easy to do so in a coherent way—that is, *ex-ante*, before the selection of portfolio weights are selected, which occurs—we perform this analysis discounting three alternative levels of transaction costs in order to capture how investors may be affected for different degrees of wealth invested and professional skill.[4] When computing optimal portfolio weights we consider three levels of risk aversion, both under MV and under power utility preferences. Finally, a variety of investment horizons, ranging from 1 to 60 months, are analyzed.

In the first part of the paper, we find that, at least in a statistical perspective (in comparisons with a simple average excess return benchmark), the three ratios display weak out-of-sample forecasting power for US equity excess returns over the sample period January 1956 – December 2012. Consistently with earlier papers (see, e.g., Almadi et al., 2014; Dangl and Halling, 2012; Lettau and Van Nieuwerburgh, 2008; Paye and Timmermann, 2006; Weigand and Irons, 2007), it turns out that the predictive power of the ratios is time-varying, as captured by a simple and yet powerful two-state MS models.[5] In the case of both linear and MS models, the strength of OOS predictability increases with the horizon.[6]

Given the modest forecasting power exhibited by the three ratios, in the second part of our paper, we assess whether the ratios are able to generate any economic value, in the form of an improvement of the realized (possibly, risk-adjusted) performance. Under simple switching

[4] Fabozzi et al. (2006) discuss the importance of considering transaction costs ex-ante. However, this is particularly tricky under dynamic portfolio strategies in the presence of regimes, see, e.g., Jang et al. (2007).
[5] The poor OOS predictive ability of our ratios has been connected to possible estimation biases due to their high persistence and to the possibility that spurious regressions issues may affect inference (see Nelson and Kim, 1993, and Stambaugh, 1999). As explained by Guidolin (2011a), MS models may capture the persistence in the ratios through the Markov chain variable that drive MS mixing and yet preserve stationarity.
[6] Samuelson (1969) and Merton (1969) showed that when investors have power utility and asset returns are identically and independently distributed over time (IID, hence unpredictable), horizon effects are absent, i.e., optimal asset allocations turn myopic. However, since the 1990s a few papers (e.g., Barberis, 2000) showed that the horizon plays a crucial role in assessing whether predictability has any economic value.

strategies, the EP ratio yields the most promising results: when we do not account for transaction costs, the switching strategy based on MS predictive regressions produces higher mean returns than the benchmark at all investment horizons. However, the superior realized performance declines with the horizon. Returns from the switching strategy are less volatile than the ones from the benchmark. However, the results are less stark when we account for transaction costs: only at long-term horizons, EP-based strategies generate incremental performances vs. the benchmarks. MS performances for the DP and BM ratios are similar but weaker, as in these cases transaction costs take a clean sweep of any risk-adjusted profits. Interestingly, we often observe that strategies that exploit MS are unable to generate higher average profits than simpler strategies and yet these may outperform in risk-adjusted terms because these imply a lower risk exposure. Yet, when transaction costs (even though these are accounted for when setting up the strategies to avoid wasteful trades) are imputed, realized returns decline enough to also compromise the risk-adjusted measures.

Next, we explicitly consider risk-averse investors. Under MV preferences, the economic value generated by all predictors turns increasingly positive when the investor becomes more risk averse. Indeed, when we consider highly risk averse investors, the MV allocation based on predictive regressions yields a higher certainty equivalent return than a strategy that disregards predictability (at least for long-term investment horizons). On the contrary, for an investor with the lowest risk aversion coefficient none of the predictors achieves better results (in terms of realized certainty equivalent) than a strategy that ignores predictability. Robustness checks extend these exercises to power utility, constant-relative risk aversion investors that use more sophisticated MSVAR models to jointly, recursively forecast both excess stock returns and predictors. Such more complex MS predictive frameworks reveal that the EP ratio always delivers some positive economic value, in terms of gains in certainty equivalent return, vs. simple no-predictability, strategies, in particular for short-term horizons. On the contrary, a strategy that exploits predictability of the BM ratio outperforms the IID benchmark only at a long-term (5-year) horizon. Finally, the gains deriving from predictability based on the DP ratio are instead modest and similar to those obtained when ignoring predictability entirely.

We obtain important insights on the economic value that accounting for elusive, unstable predictability (see Timmermann, 2008)—here captured in a MS framework—may offer to asset managers. On the one hand, and consistently with the earlier literature, capturing shifts in predictability regimes moderately improves OOS statistical performance and offers a chance of creating economic value to investors. On the other hand, this occurs mainly in favor of long-horizon investors who are highly risk averse. This occurs because MS models represent poor devices to implement market timing strategies, i.e., the switching dynamics may be modelled but never precisely predicted. Equivalently, MS models are in general mediocre tools to forecast the short-run dynamics of the conditional density of excess returns. On the contrary, MS models offer accurate predictions of the shape of the long-run, essentially unconditional density of excess returns. In this respect, MS models provide more precise estimates of the shape (thickness) and asymmetries (skewness) of such long-run densities and hence allow risk-averse investors to appropriately hedge their portfolios thus obtaining stronger risk-adjusted performances. In fact, this occurs even though accounting for regimes tends to reduce average realized performances, especially net of transaction costs.[7]

Finally, we also perform robustness checks using relatively sophisticated MSVAR models in which the joint dynamics of excess equity returns and the predictors is captured and we assume preferences that integrate over the entire (conditional) predictive density to allow for expected utility maximization.[8] Such additional back-tests fail to lead to any marked improvement in ex-post realized performance. In fact, often the strategies that imply frequent rebalancing (and this is anticipated as a possibility already at inception in the choice of weights) lead to inferior realized performances. Although additional tests (e.g., with other predictors) and more sophisticated regime switching models (e.g., threshold models or MS with time-varying transition probabilities) may yield better outcomes, it seems that most of the benefits of MS can be harvested using relatively simple statistical approaches and that these mostly concern long-term, highly risk averse investors.

The rest of the paper is structured as follows. In Section 2 we introduce our data set. Section 3 provides details on the statistical in-sample and OOS results. In Section 4, we describe the investment strategies and report our key findings on the economic value of financial ratio-driven

---

[7] This finding echoes Lettau and van Nieuwerburgh's (2008) result that it is the uncertainty in the estimation of the nature and size of the steady state shifts in predictive relationships rather than the estimation of their exact dates that is ultimately responsible for the failure of the real-time out-of-sample linear predictions to beat the benchmark average sample return.

[8] This extension follows Guidolin and Timmermann (2007) who investigate the effects of regimes in stock and bond returns for asset allocation under constant relative risk aversion preferences. They find that even when portfolio weights depend on regimes, the dividend yield remains a relevant predictor. Our paper extends their results to different predictors, alternative preferences, and - at least in the simplest framework - accounts for transaction costs.

predictability. In Section 5, we test the robustness of our findings when preferences feature constant relative risk aversion and under dynamic MSVAR models. Section 6 concludes.

2. The Data

In order to assess the predictive power of alternative ratios we base our research design on the same data as Welch and Goyal (2008). We collect monthly US data on equity market returns, risk free rate, dividend, earnings, and book-to-market ratios spanning the January 1926 - December 2012 sample. Our dependent variable, i.e., excess equity returns, is calculated as the difference between returns on the Standard & Poor Index, provided by the Center for Research in Security Press (CRSP) and the risk-free rate, approximated by the US Treasury bill rate. As for the latter, we use the "3-month Treasury Bill Secondary Market rate" series from the Federal Research Economic Data (FRED) repository of the Federal Reserve Bank of Saint Louis for the period 1933-2012. Because this series is not available before 1933, we use the "US Yields on Short-Term United States Securities, Three-Six Month Treasury Notes and Certificates, Three-month Treasury" series, available from the National Bureau of Economic Research (NBER) database for the earlier sub-sample, 1926-1933.

The first predictor, the dividend-price ratio, is computed as the natural logarithm of the ratio between monthly dividends and price: at the end of each month in the sample, we take the 11-month lagged value of the S&P 500 index as the denominator and the 12-month moving average of the dividends paid by all the shares included in the index as the numerator. Similarly, the second predictor, the earnings-price ratio, is obtained as the natural logarithm of the ratio between the 12-month moving average of the earnings reported by all S&P 500 companies and the 11-month lagged value of the index. Finally, the third predictor, the book-to-market ratio, is calculated as the ratio between the aggregate book value of all Dow Jones Industrial Average companies and the market capitalization of the index, as in Welch and Goyal (2008).[9] We focus on these three predictors because the bulk of both the academic and practioners' literatures have often returned on their business cycle properties and predictive performance (see e.g., Campbell and Shiller, 2001; Cardinale et al., 2014, Weigand and Irons, 2007, and references therein).

[9] The data on the earnings and the dividends of the S&P 500 companies are from Robert Shiller's website for the period 1926-1987 (at http://www.econ.yale.edu/~shiller/data.htm). For the period from 1988 to the end of the sample, we use the S&P Corporation dataset. For the second part of the sample we perform the interpolation of quarterly earnings suggested by Welch and Goyal (2008). Finally, the book values come from the Value Line's website, in particular their *Long-Term Perspective Chart* of the Dow Jones Industrial Average.

Table 1 presents the key summary statistics for stock returns, excess equity returns, and the three predictors. The average monthly return is 0.80%, implying an average annual return of 9.6%, while the annualized volatility is 19%. The return in excess of the risk-free rate is 0.50% per month (6% on an annualized basis) and corresponds to a Sharpe ratio of approximately 0.32. All the series show deviations from a normal distribution and we can always reject the null hypothesis using a Jarque-Bera test based on deviations of skewness and kurtosis from the Gaussian values of zero and three, respectively. In particular, the distribution of equity returns and excess returns is left-skewed and displays fatter tails than a normal distribution (see Jones and Stine, 2010; Stoyanov et al., 2011).

3. Statistical Evidence of Predictive Ability of Alternative Ratios

In this section, we develop a set of recursive out-of-sample (OOS) exercises to assess the (pseudo) real time predictive power of the three valuation ratios introduced in Section 2, namely the DP, the EP, and the BM ratios. We rely on two alternative statistical frameworks: standard, single-state univariate regressions, one for each predictor; Markov switching regressions that capture any instability in the predictive relationships in the data.

3.1 Univariate linear regressions

Following Goyal and Welch (2008), we estimate a set of standard, single-state univariate regressions, of the form:

$$r_{t+h} = \alpha_h + \beta_h X_{\mathrm{t}} + \varepsilon_{t+h}\,, \tag{1}$$

where $r_{\mathrm{t+h}}$ is the excess return at time t+h, α_h is the intercept, β_h is a slope coefficient, X_{t} is the predictor at time t (DP_{t}, EP_{t}, and BM_{t}, respectively), and $\varepsilon_{\mathrm{t+h}} \sim N(0, \sigma^2_{t+h})$ is the residual. The forecast horizon h is set equal to 1, 6, and 60 months, respectively, so that we obtain three regressions for each selection of a predictor. The results of the estimation of these regressions are reported in the first column of Table 3. We observe that all the proposed ratios turn out to be statistically significant in-sample predictors of excess returns when we set the forecasting horizon equal to 1 or 6 months. Instead, the null of β_h equaling zero can never be rejected when we set h to 60 months. So, in a statistical perspective, there is in-sample evidence of predictability over short horizons but not much predictive power persists at the longest horizons.

In order to assess and compare the OOS predictive ability of the three ratios, we estimate the forecasting regressions recursively on the basis of the data available only up to a given time t, where

t expands over a period ranging from January 1951 to December 2012.[10] Figure 1 plots the evolution over time of the slope coefficients associated to the three predictors at a 1-month horizon. The coefficient of the DP ratio is rather stable over the whole period, with a small estimated value of approximately 0.02. Instead, the estimated coefficients of the EP and BM ratios visibly decline over time, from 0.03 to 0.01 and from 0.04 to 0.02, respectively. To evaluate the OOS forecasting power of our predictors we compute the RMSFE (root-mean-square forecast error), defined as follows:

$$RMSFE(h) = \sqrt{\frac{\sum_{i=1}^{n}(\hat{r}_{i,t+h} - r_{i,t+h})^2}{n}}, \tag{2}$$

where $\hat{r}_{i,t+h}$ is the horizon-h forecast of the excess return, $r_{i,t+h}$ is the realized excess return, and n is the number of times that the recursive exercise is repeated. We also compute the RMSFE of a no-predictability Gaussian IID benchmark with constant mean and variance and compare it with the one implied by the univariate predictive regressions. In practice, we subtract the RMSFE of each predictive regression to the one of this baseline model, such that a negative difference means that a predictive model implies a worse predictive accuracy than the benchmark, while a positive difference signals that our model is more accurate than the simple IID model that hinges forecasts on a recursively estimated sample mean. We perform the exercise for each of our predictors at the three prediction horizons listed above.

Figure 2a plots the differences in RMSFE between the baseline model and a regression on the DP ratio. Interestingly, at short horizons (1 and 6 months), the two models yield approximately the same accuracy. Instead, at a 5-year horizon, the performance changes abruptly over time: before the 1970s, the naïve, recursive sample mean displays stronger predictive power than the dividend-price ratio, while the situation flips around later in the sample. Figure 2b shows the differences in RMSFE between the baseline model and a regression on the EP ratio. The results are similar to the ones obtained for DP: the two models show comparable forecasting performances at 1- and 6-month horizons; instead, at a 5-year horizon the EP ratio yields a slightly higher accuracy. Finally, Figure 2c compares the baseline model with a regression on the BM ratio. In this case, the "winning model" changes over time: until the 1970s the BM ratio outperformed the unconditional mean at the 1- and 6-month horizons, while it underperforms in the subsequent period; however, exactly the opposite happens when we consider a 5-year horizon. In contrast with the in-sample evidence,

[10] The initial coefficient estimates are obtained using data from January 1926 to December 1950. This implies that our first forecast regression predicts the excess return for January 1951 (h=1 month), for June 1951 (h= 6 months), and for December 1955 in the case of h=60 months. For the sake of homogeneity and to avoid showing results based on few observations only, all our figures in what follows start out in January 1966.

our OOS results shows that genuine forecast power actually *increases* with the return horizon. This is consistent with findings reported in previous literature (see, e.g., Fama and French, 1988, who report that the R^2 statistic increases strongly with the horizon when the predictor variable is highly persistent, as our valuation ratios are and more recently Almadi et al., 2014).

3.2 Univariate Markov switching regressions

In Section 3.1, we have noted that none of our predictors was able to persistently outperform a baseline model where we forecast returns employing their unconditional, recursive sample mean. In addition, we have emphasized some degree of instability in the predictive ability of the three ratios. Consequently, in this section, we extend the empirical evidence by using a simple regime switching process driven by a standard first-order, homogeneous, irreducible and ergodic Markov chain to accommodate the potential time-varying predictive power of our predictors. In particular, we specify the following Markov switching intercept heteroskedasticity (henceforth, MSIH), model:

$$r_{t+h} = \alpha_{S_{t+1},h} + \beta_{S_{t+1},h} X_{\mathrm{t}} + \sigma_{S_{t+1}} \varepsilon_{t+h} \qquad \epsilon_{t+h} \sim \text{IID } N(0,1), \tag{3}$$

where the intercept, slope coefficients, and the variance all depend on an unobservable state variable S_t. Regime switches in S_t are assumed to be governed by the transition probability matrix, **P**, with elements:

$$\Pr(s_t = i | S_{t-1} = j) = p_{j,i}, \quad i,j = 1, \ldots, k \tag{4}$$

where k is the number of regimes assumed in the analysis. The parameters are estimated by optimizing the likelihood function associated with (3) and (4). Because the state variable is unobservable, this should be treated as a latent variable and thus we use the EM algorithm to update our parameter estimates, as suggested by Hamilton (1989).[11] Importantly, the EM algorithm will naturally deliver time-varying (smoothed) estimates of state probabilities of the "system" defined by the econometric model to be in each of the regimes at each point in time.

To determine the number of regimes we conduct an in-depth specification search with values of k = 1, 2, 3. Notably, when k = 1, our specification collapses back to the linear model described in equation (1). To assess if a number of regimes higher than one is required we perform the test proposed by Davies (1977) and we find that the single-state linear model is always misspecified. Table 2 reports the values of the Hannan-Quinn (HQ), Akaike (AIC), and Bayes-Schwartz (BIC)

[11] For an in-depth review of the estimation methods required to perform the required model calibrations, please refer to Guidolin (2011a) or the textbook introduction in Fabozzi et al. (2006).

information criteria, which trade off fit against parsimony, for alternative specifications of each model.[12] Both the HQ criterion and the AIC point towards a three-regime model. However, for all the predictors and the forecasting horizons the more parsimonious BIC indicates a two regime model. Accordingly, and to offer the simplest, starkest case, in what follows we recursively estimate a MSIH (2) model.

3.3 Model estimates

Table 3 reports parameter estimates for the two-state MS regressions for each of the three predictors at 1-, 6-, and 60-month horizons. It is easy to provide a common interpretation of the two regimes across the different regressions: regime 1 is always characterized by high volatility (on average, approximately three times the volatility in regime 2) and a shorter duration than regime 2. Regime 2, instead, represents a low volatility, very persistent regime (approximately 8 years of average duration vs. 12 months in the case of regime 1). As a result, we shall refer to regime 1 as the high-volatility regime and to regime 2 as the low-volatility regime.

The first panel of the table reports parameter estimates for three regressions of excess stock returns on the DP ratio. While this ratio shows substantial (in terms of the size of the estimated coefficients) predictive power for excess returns at all forecasting horizons in the low-volatility regime, the corresponding β coefficient is only significant at the 1-month horizon. The second panel contains parameter estimates for the EP ratio regressions. Interestingly, in contrast to the linear model, where we were able to find a statistical significant predictive relationship between EP and excess returns, in the MS case we are never able to reject the null hypothesis that β is equal to zero at a 10% size. Finally, the last panel of the table reports the results for the BM ratio. Contrary to DP, the BM ratio retains some forecasting power only in the high-volatility regime at both 1- and 6-month horizons. This finding suggests that drawing conclusions as to which ratio is able to better forecast excess equity returns may heavily depend on the market regime assumed by an analyst (Almadi et al., 2014; Dangl and Halling, 2012; Gonzalo and Pitarakis, 2012). All in all, even when regimes are taken into account, the evidence of in-sample predictability remains mixed and it tends to concentrate in different regimes for different predictors. Moreover, also in the case of MS

[12] As discussed in Brooks (2014) among other textbook treatments, the goal of a researcher employing information criteria in model selection consists in minimizing them.

predictive regressions, across different regimes and predictors, the longer the horizon is, the weaker the strength of the in-sample evidence of predictability.

3.4 The forecasting performance of Markov switching models

In this section we perform the same OOS exercise discussed in Section 3.1, now applied to the MS regressions estimated in Section 3.3. This allows us to both compare the predictive ability of the three ratios and to assess whether a model that accounts for regime shifts may deliver better predictive performance than a simple linear regression. To provide an example of how an asset manager will operationalize the predictability framework offered by a MS model such as the ones in Table 2, Figure 1 plots the recursively estimated β coefficients in the high-volatility and in the low-volatility regimes at a 1-month horizon. In addition, the dashed line represents the evolution over time of the smoothed β computed as the weighted average of regime-specific slope coefficients with weights equal to the smoothed probabilities. Notably, in the high-volatility regime the β coefficients of all the three predictors are visibly higher than in the low-volatility regime, thus implying that the predictive power of the different ratios notably changes over time.

Figures 2a through 2c compare the predictive accuracy of the MS regressions both with a Gaussian IID model and the standard linear regressions discussed in Section 3.1. Figure 2a plots the dynamics of the differences in RMSFEs implied by the DP ratio at different horizons. In particular, a positive difference signals that the two-state MS regression outperforms the benchmark model. Interestingly, the MS regression that uses the DP ratio consistently outperforms both the Gaussian, sample-mean model and the linear regression at a 6-month horizon; however, it underperforms both benchmarks at a 1-month horizon. The finding is less clear at a long-term (5-year) horizon, as the MS model underperforms its benchmarks until the 1970s, while the opposite holds later.

Figure 2b performs the same comparison for the EP ratio. In this case, the MS model is not able to significantly outperform either the standard sample average model or the linear one. Conversely, it underperforms both benchmarks at a 1-month time horizon in the period from 1972 to the end of the sample and at a 5-year horizon until the 1970s. Finally, Figure 2c shows our OOS RMSFE findings concerning for the BM ratio. Also in this case, the MS model is in general unable to deliver better predictions than a simple Gaussian model or a linear regression, particularly at 1- and 6-month horizon. In conclusion, both linear and MS predictive regressions occasionally outperform a simple IID model that uses the unconditional mean to forecast aggregate excess stock returns.

However, the statistical evidence of predictive ability of the three ratio is mostly weak and never persistent over time, especially for short-term horizons. However, as discussed in the Introduction, because an investor ought to be chiefly interested in the economic value that using the financial ratios as predictors may deliver, in the next section we exploit the forecasts from all of our models to build simple and therefore insightful investment strategies that exploit any predictability in the data, however weak.

4. Investment strategies

4.1 Switching strategies

In this section, we assess whether a simple switching investment strategy *à la* Pesaran and Timmerman (1995), which allocates all the available wealth alternatively in stocks or Treasury bills (which are assumed to yield a risk-free rate), is able to outperform a passive buy-and-hold (B&H) strategy on the S&P index when the former is based on the statistical evidence of predictability reported in Section 3. In practice, we build our switching strategy such that at time t the investor decides her asset allocation until $t+h$ depending on the sign of her forecast of the excess equity return up to time $t+h$. If the model predicts a positive excess outperformance of the stock index, the investor allocates all her wealth to equities, otherwise she simply invests in Treasury bills. The exercise is recursively repeated for a sample spanning the period 1956 - 2012, so that on every month the investor selects her optimal portfolio based on all the data available up to that point.[13] We consider three different holding periods (1, 6, and 60 months) matching the forecasting horizons already examined in Section 3. When the investment period exceeds 1 month, we implicitly consider "overlapping" investors (i.e., the first investor decides her allocation from t to $t+h$, the second one invests between $t+1$ and $t+h+1$ and so on). The realized value of wealth at the end of each holding period is equal to:

$$W_{t+h} = W_t[(1 + R_{t+h})I(ER)_t + (1 + Rf_{t+h})(1 - I(ER)_t], \tag{5}$$

where W_t is wealth at the beginning of the period, R_{t+h} is the equity return between t and $t + h$, Rf_{t+h} is the risk-free rate between t and $t + h$, and $I(ER)_t$ is a dummy variable that equals 1 when all wealth is invested in the index and 0 otherwise. The average realized performance of the

[13] Our initial regression is based on data from December 1955 to December 2012, such that the first predicted excess return concerns January 1956 for h=1 month, June 1956 for h= 6 months and December 1960 for h=60 months. Similarly to Section 3, for the sake of comparability and to avoid showing results based on few observations all our charts start from January 1966. This also applies to all the exercises presented in Sections 4 and 5. Such expanding window schemes are now typical in the literature, see, e.g., Almandi et al. (2014).

recursive implementation of this strategy is compared with the average performance of the B&H strategy where all the initial wealth is invested in the equity index until the end of the sample period. We compute (monthly) realized mean, variance, and the Sharpe ratio (see Hubner, 2007, for a justification, as the resulting managed portfolio is an exclusive, directional investment vehicle) of returns from both strategies to assess whether the use of the forecasts generated by our predictive regressions may produce better results than a passive strategy, both in terms of improving average realized returns or reducing their variance (thus increasing risk-adjusted performance). The exercise is performed twice, using the forecasts from single-state as well as two-state MS models to investigate whether modelling and predicting regime changes leads to an increase in the economic performance of the predictive regressions.

When we backtest investors' decisions in real time, it is of paramount importance to account for transaction costs, especially considering that our benchmark is a zero cost (passive) B&H strategy that in reality can be easily replicated with a cheap index fund. In the case of the active strategies introduced above, an investor will incur costs any time she decides to switch from equity to risk-free bonds and vice versa. In particular, any time she trades, she will incur costs related to bid-ask spreads, stamp duties, and commissions. The estimation of the level of such costs is of course not trivial. Following Bhardway and Brooks (1992), Lesmond et al. (1999), and Stoll and Whaley (1983), Lynch and Balduzzi (2000) report and impute round-trip costs of 100 bps for investor trading directly individual stocks on the NASDAQ and the NYSE. Keim and Madhavan (1997) show that for large and liquid stocks, institutional investors pay a round-trip cost of approximately 38 bps. However, as pointed out by Lynch and Balduzzi (2000), the cost of trading future contracts on the S&P 500 index is much lower and may vary from 2 to 5 bps (although this estimate does not take into account margin requirements and rolling costs caused by the expiry of the front-end future in case of long investment horizons). In our analysis, we present the results for three different levels of (proportional) transaction costs: zero, 25 bps and 50 bps. The investor will consider the transaction costs when deciding to switch from the risk-free asset to equity: if the cost of switching is higher than the predicted equity excess return she will not change her allocation; similarly, if the negative excess return from staying in stocks is less than the transaction costs, she will refrain from closing her position in the index.

The realized performances of the alternative strategies at different horizons are reported in Table 4. The first panel shows the results for the DP ratio. Interestingly, the switching strategy is never able to produce mean returns higher than the B&H strategy, irrespective of the holding period, even under zero transaction costs. Indeed, the B&H strategy yields a mean return of 0.77% per month, while a switching strategy based on linear forecasts generates returns of 0.73% and 0.66% per month for short and long investment horizons, respectively (again, assuming zero transaction costs). These results do not improve much even when we use the MS model to forecast S&P 500 index returns. However, the returns of the switching strategy become less volatile than the ones of the B&H strategy. This is not surprising because we verify that the switching strategy tends to suggest to invest in the less volatile Treasury bills at any time when the forecasts for equity returns are negative, i.e., during the high volatility regime. Occasionally this realized reduction in volatility more than compensates the lower average realized returns from the switching strategy leading to an outperformance of the latter in terms of realized Sharpe ratios. In particular, a linear predictive regression shows higher risk-adjusted performance than the B&H strategy at short-term investment horizons (but not at long-term ones) when we consider zero transaction costs (the Sharpe ratio is 0.089 vs. 0.084). Yet, this advantage disappears when we consider transaction costs. Indeed, under high transaction costs, the Sharpe ratio of a MS-based strategy at short-term investment horizons even turns negative as the realized returns do not exceed the risk-free rate.

The second panel shows the results for the EP ratio. This ratio yields a higher predictive power (in terms of ability to generate positive economic value) than the DP ratio. Indeed, if we do not account for transaction costs, the switching strategy based on MS predictive regressions produces higher mean returns than the benchmark at all investment horizons. The superior realized performance equals 0.05% per month for short-term horizons, but it is only 0.02% for long-horizon ones. In addition, also in this case, the returns from the switching strategy are less volatile than the ones from the benchmark, so that the stronger performance is more evident in terms of risk-adjusted returns. However, the situation is less clear when we account for the transaction costs. Indeed, if we consider short-term investment horizons, the switching strategy is never able to outperform the benchmark either in terms of realized returns or of risk-adjusted criteria. However, for long-term (5-year) horizons, the strategy based on linear predictive regressions generates a return of 0.82% per month vs. 0.77% for the B&H benchmark.

Finally, the last panel of Table 4 contains the results concerning the economic value assessment for the BM ratio. Also this ratio seems to generate value when compared to the passive strategy when we ignore transaction costs, especially at short investment horizons. Indeed, a switching strategy based on MS forecasts leads to realized average returns equal to 0.80% (at a 1-month horizon) and 0.78% (at a 6-month horizon) to be contrasted to the 0.76% of the B&H strategy and to Sharpe ratios of 0.975 and 0.942, respectively (versus 0.837 from the B&H benchmark). Also for long term investment horizons, the MS strategy generates a higher risk-adjusted performance, even if it does not beat the benchmark in terms of raw average returns. However, as soon as transaction costs are included, the dynamic strategies lose their uphand vs. the passive one.

In conclusion, the use of the DP ratio as a predictor of equity returns does not seem to generate better performances than a simple B&H passive strategy, even if we ignore transaction costs. Notably, this is quite surprising if we consider the *hit-ratio* of our MS predictive regression, i.e., the percentage of correct signs of excess return predictions. Indeed, a MS predictive regression based on DP succeeds in forecasting the sign of stock excess returns 49.5% of times for 1-month investment horizon, 51.4% of times for 6-month investment horizon and 67.5% of times for 60-month investment horizon (vs. 41.3%, 49.80%, and 59.90% of a linear predictive regression). The *hit*-ratio of MS predictive regressions is even better in the case of EP and BM ratios, where it exceeds 70% for long-term investment horizons. Considering that the strong ability to correctly predict the signs of the excess returns do not turn into large economic gains, this may indicate that MS predictive regressions fail to forecast large negative returns (see e.g., Famy, 2007; Leung et al., 2000). However, we should also keep in mind that, while a strong performance in terms of correct sign predictions has been previously noticed to be necessary for large economic value to be generated (see e.g., Leitch and Tanner, 1991), this does not represent a sufficient condition.

Yet, when we forecast returns using the BM and, especially, the EP ratio, we are able to produce some economic value, at least under zero transaction costs (which may indeed represent a realistic assumption in the case of some large institutional investors). In addition, if we consider their risk-adjusted performance (measured by the Sharpe ratio), switching strategies based on EP and BM ratios occasionally outperform the benchmark even when we deduct rather massive transaction costs (especially for long-term horizons). Interestingly, DP-driven forecasts – arguably the ones who have received the most attention in the literature (see e.g., Barberis, 2000, and references therein) – are the only ones that fail to support economic value generation on a robust scale. Finally, we notice that MS predictive regressions generally deliver similar (and occasionally better) realized returns than the linear ones, with one single exception: in the case of long-term horizons simple linear regressions perform better than the MS ones.

4.2 Mean-variance strategies

In this section, we follow the approach proposed by Campbell and Thomson (2008) to assess whether predictive regressions alternatively based on DP, EP, or BM ratios are able to generate a meaningful improvement in realized portfolio performance to an investor characterized by a standard mean-variance utility function defined over terminal wealth, W_{t+h}, with a coefficient of risk aversion $\gamma > 1$, and an investment horizon of h months:

$$U(W_{t+h}) = E[W_{t+h}] - \frac{\gamma}{2} Var[W_{t+h}]. \quad (6)$$

The investor is supposed to maximize the utility function by selecting at time t (and holding such a selection until $t+h$) the optimal allocation ω^* to equity, while $1 - \omega^*$ is invested in risk-free Treasury bills. In order to prevent short-selling, we constrain ω^* to be positive and lower than or equal to 100%. Simple algebra (see Campbell and Thomson, 2008) reveals that the optimal conditional portfolio rule is

$$\omega^* = \left(\frac{1}{\gamma}\right) \frac{\mu + x_t}{\sigma_{\varepsilon,t}}, \quad (7)$$

where x_t is a function of our predictor X_t, namely $x_t = \hat{\alpha} + \beta X_t - \mu$. Notably, an investor's portfolio allocation at time t depends on her forecast of r_{t+h}, which is based on the observed value of our predictor. Consequently, the denominator of the formula is the standard error of the residual of our predictive regressions.

After having determined the optimal recursive weights using (7), we also compute the ex-post, realized certainty equivalent value (CEV), and compare it with the matching CEV from a MV strategy based on a simple IID model where the best forecast of future returns is the sample mean. The exercise is recursively repeated from 1951 to 2012, following the same logic already adopted for the switching strategy discussed above (i.e., we consider overlapping investors for holding periods longer than one month). Table 5 shows the average differences in the realized utility (Δ CEV) of the mean-variance strategies alternatively based on our predictive regressions and on an IID model, at 1-, 6-, and 60-month investment horizons and for different values of the risk aversion coefficient (γ=2, 5, and 10). A positive value of this difference implies that a MV strategy based on the predictive regressions outperforms the one based on an IID model in terms of risk-adjusted realized portfolio performances. The results for each predictor are based on our two alternative predictive models, linear vs. MS regressions. In addition, we report the test statistics of a difference-in-means test, to assess whether ΔCEV is significantly different from zero. In particular, the test statistic is computed as follows:

$$t = \frac{\Delta\overline{CEV}}{\sqrt{(\hat{\sigma}^2{}_X + \hat{\sigma}^2{}_{IID} + \hat{\sigma}_{X,IID})/n}}, \tag{8}$$

where n is the number of observations in our pseudo OOS period, $\hat{\sigma}^2{}_X$ is the sample variance of the realized returns from the MV strategy based on a predictive regression, $\hat{\sigma}^2{}_{IID}$ is the sample variance of the realized returns from the MV strategy based on an IID model, and $\hat{\sigma}_{X,IID}$ is the sample covariance of the two sets of realized strategy returns.

In general terms, the most salient result that we can detect is that the economic value generated by predictability turns increasingly positive when the investor becomes increasingly risk averse. Indeed, when we consider γ=10 (which represents extremely risk averse individuals) the MV allocation based on predictive regressions yields a higher CEV than a strategy that disregards predictability (at least for long-term investment horizons). On the contrary, for an investor with the lowest risk aversion coefficient (γ=2) none of our predictors seems to achieve better results (in terms of realized CEV) than a strategy that ignores predictability. Indeed, at long (5-year) investment horizons, the MV strategy based on our three predictors even underperforms the benchmark, thus generating a negative and statistically significant ΔCEV.

The first panel of Table 5 reports the results for the DP ratio. The ability of predictability to generate meaningful economic value is rather modest. Indeed, ΔCEV is positive and statistically significant only when we consider γ=10 and a long-term (5-year) investment horizon. In addition, only a regime-switching predictive regression delivers a positive utility gain (ΔCEV equal to 0.0501 vs. -0.0624 of the linear predictive regression). The second panel shows the results for the EP ratio, and the third one for the BM ratio. The results are very similar to the ones already commented for DP: in particular, only a MS regression is able to outperform the simple IID model, but this is limited to the case of a strongly risk-averse investor with a long-term holding period (ΔCEV equals 0.0508 and 0.0557, for EP and BM ratio, respectively).

Finally, Figure 3 plots the OOS realized ΔCEV computed between the strategy based on MS predictability and the one based on a simple IID model, at different investment horizons and for different levels of risk aversion. MS predictability is never able to produce a significant utility gain at a short-term investment horizon, no matter the predictor we decide to use. However, it is possible to exploit predictability over longer time windows (e.g., 5-year, and, occasionally 6-month horizons) when the investor is highly risk-averse. Interestingly, although the strategy based on MS predictability at 5-year horizon outperforms the benchmark for a large part of the sample period, it has a poor performance during the early years 1966-1969, for all the predictors. Indeed, that period was characterized by returns that were much lower than the predicted ones, meaning that all models and especially MS did over-forecast excess stock returns.

In conclusion, the results of this exercise are different from the ones obtained from simple switching strategies in two ways. First, when a MV portfolio choice algorithm is recursively implemented, the performance of different ratios is much more homogeneous, to the extent that it becomes now difficult to identify a predictor that clearly outperforms the others using a risk-adjusted portfolio metric. Secondly, a MV strategy that accounts for predictability delivers better results at long-term horizons, while the switching strategy outperforms the benchmark at short horizons. In addition, while in the case of switching strategies the MS regressions were underperforming a linear model for longer holding periods, the reverse is true when we implement MV allocation. These findings must be related to the power of MS models to better measure and predict the dynamics of risk over time, in such a way that when optimal allocations are computed taking ex-ante risk into account (as a MV investor obviously does) and realized performances are ex-post estimated in a risk-adjusted manner, it takes medium-to-long investment horizons for MS to outperform the benchmarks, consistently with the risk management and portfolio evidence reported by Guidolin and Timmermann (2006, 2007).

5. An Extended Framework: Dynamic Portfolio Strategies in Dynamic Models

In Section 4, we have analyzed the implications of predictability, in terms of portfolio value generation, for MV investors with different degrees of risk aversion. Yet, a one may argue that a MV investor only takes into account the first two moments of the distribution of excess returns, which is in contrast with the evidence of the presence of skewness and kurtosis arising in Table 1. This also disregards the much advertised ability of mixtures of normal distributions – of which our MS predictive framework represents one such case – to capture non-normalities in the predictive density of excess risky returns (see Guidolin, 2011a, for a discussion and references). Consequently, in this section we test the robustness of our results by extending our analysis to an investor characterized by a power utility function over terminal wealth, a coefficient of risk aversion $\gamma > 1$, and an investment horizon of h months:

$$U(W_{t+h}) = \frac{W_{t+h}^{1-\gamma}}{1-\gamma}. \tag{9}$$

It is well known (see Bekaert et al., 1998; Jones and Stine, 2010; Malevergne and Sornette, 2005) that besides implying a constant coefficient of relative risk aversion, power utility preferences make expected utility of an investor a function of predictive skewness, predictive excess kurtosis, and more generally of the relevant and possibly time-varying features of predictive density functions of terminal wealth. We follow a recent literature (see e.g., Guidolin and Timmermann 2007; Guidolin and Hyde, 2012) and investigate whether more realistic preferences and robust models that endogenize returns and predictors may reveal different results on the value of MS predictability.

In parallel to the comparison we have performed between single- and two-state predictive regressions, we extend the earlier econometric frameworks by adopting two alternative models to forecast excess returns. First, following Barberis (2000), we use a standard VAR (1) model:

$$\binom{r_t}{x_t} = \binom{\mu}{\mu_x} + A\binom{r_{t-1}}{x_{t-1}} + \binom{\varepsilon_t}{\varepsilon_{x,t}}, \tag{10}$$

where r is the excess return, x is the predictor, and $(\varepsilon_t, \varepsilon_{x,t}) \sim N(0, \Omega)$. This single-state model generalizes our earlier predictive regression to the case in which both returns and the predictor are endogenous. Second, following Guidolin and Timmerman (2007), we extend our analysis to an investor that forecasts excess returns using a Markov switching VAR, which allows a VAR-type relationship between the predictor and equity returns to vary across different states. In particular, we perform a standard model selection process to select a two-state VAR (1).[14]

Similarly to Barberis (2000) and Guidolin and Timmerman (2007), we consider both a framework where the initial allocation is kept fixed till the end of the holding period (which is set equal to 1, 6, and 60 months) and an allocation problem where the investor is able to rebalance her portfolio dynamically at regular intervals. In practice, the terminal wealth of the investor equals:

$$W_{t+h} = (1-\omega)e^{Rf_{t+h}} + \omega e^{R_{t+h}}, \tag{10}$$

where the initial wealth is normalized to 1, h is the investment period, ω is the portion of initial wealth invested in equity, Rf_{t+h} is the risk free rate that applies between t and $t+h$, and R_{t+h} is the

[14] The results of the specification search and the estimates of the VAR (1) and MSVAR (2,1) (i.e., characterized by two regimes and one lag) models are not reported to save space but are available upon request.

equity return over the holding period. The investor maximizes the expected utility of her terminal wealth conditional on the information available at time t:

$$max_{\omega} E_t\left(\frac{\left[(1-\omega)e^{Rf_{t+h}}+\omega e^{R_{t+h}}\right]^{1-\gamma}}{1-\gamma}\right), \tag{11}$$

where ω is constrained such that $\omega \in [0,1]$ to avoid short-selling. Consequently, to find the optimal weights the investor has to solve:

$$max_{\omega} \int \frac{\left[(1-\omega)e^{Rf_{t+h}}+\omega e^{R_{t+h}}\right]^{1-\gamma}}{1-\gamma} p\left(R_{t+h}\middle|z,\hat{\theta}\right) dR_{t+h} \tag{12}$$

where $p\left(R_{t+h}\middle|z,\hat{\theta}\right)$ is the distribution of future returns conditional on the estimated parameters $\hat{\theta}$ and on past values of returns. We approximate the integral for expected utility with Monte Carlo techniques by drawing a number M equal to 20,000 of simulated R_{t+h} from our return generating process and computing:

$$\frac{1}{M}\sum_{i=1}^{M}\left(\frac{\left[(1-\omega)e^{Rf_{t+h}}+\omega e^{R_{t+h}}\right]^{1-\gamma}}{1-\gamma}\right). \tag{13}$$

We perform the computation for ω=0.01, 0.02, ..., 0.99 and choose the weight that maximizes (13).

When we also consider the possibility that the investor optimally rebalances her portfolio over the holding period, then we have to solve a dynamic programming problem. We divide the investment horizon into J intervals $[t_0, t_1], \dots, [t_{k-1}, t_j]$ such that the investor rebalances her portfolio at time $(t_0, t_1, \dots, t_{j-1})$. In our case each interval is equal to one year; consequently, the results of rebalancing and non-rebalancing strategies will be identical whenever the holding period is lower than one year. At this point we solve the problem by using backward induction.[15]

Figure 4 shows the average compensatory fees which make an investor with a risk aversion coefficient γ = 5 and holding periods of 1, 6, and 60 months indifferent between a strategy that exploits predictability and one that ignores it. We present the results for both the MSVAR and the linear VAR model for an investor that rebalances her portfolio within the holding period and a one who does not. Interestingly, when we use a MSVAR to model predictability, the EP ratio is always able to deliver some positive economic value, in terms of gains in certainty equivalent, with respect to a simple IID strategy, in particular for short-term horizons. On the contrary, a strategy that exploits predictability of the BM ratio outperforms the IID benchmark only at a long-term (5-year) horizon. Finally, the gains deriving from predictability based on the DP ratio are modest and limited to the case of a linear strategy without rebalancing with a 5-year investment horizon.

[15] For discussion of the solution to the dynamic programming problem see Guidolin and Timmerman (2007).

Interestingly, we notice that the compensatory fees are much lower (and often negative) for an investor who is allowed to rebalance her portfolio with yearly frequency than for a buy-and-hold investor. Indeed, it seems that an investor who does not have the chance to rebalance her portfolio manages to achieve higher gains from predictability. This is especially true as far as linear predictability is considered. However, this is not surprising: in fact, a linear model is misspecified because it disregards the existence of regimes; consequently, when we iterate it to forecast long-run densities and hence compute optimal long-run strategies in the case of rebalancing strategies, the errors cumulate over-time, leading to negative compensatory fees. All in all, these findings confirm the robustness of our earlier results that improvements in realized portfolio performances are possible by exploiting predictability, especially when MS are taken into account and especially for long-horizon investors. The dominance of EP over the other predictors and especially DP also emerges to be stark, when risk-averse individuals are considered.

6. Conclusions

This paper has adopted both statistical and economic approaches to investigate whether standard (aggregate) valuation ratios (namely, the dividend-price, the earning-price, and book-to-market ratios) show evidence of predictive power for excess stock returns that may be concretely exploited by asset managers to achieve a risk-adjusted outperformance. We contribute to the long-standing literature on predictability of excess stock returns at least in two ways. First, we propose an extensive set of different strategies that are popular among asset managers—from simple switching strategies that invest alternatively in stocks or the risk-free asset, to dynamic portfolio allocations—in order to assess whether increasing sophistication leads to superior ex-post realized performance. We find that the implementation of more sophisticated frameworks that consider also higher moments of excess returns (i.e., skewness and kurtosis) does not strengthen the allocation strategies that exploit predictability. In particular, when the investor is let free to rebalance her portfolio with annual frequency, predictability struggles to generate any additional economic value. Second, we recognize that most predictive relationships may be time-varying and that the use of regime-switching models to capture such instability may lead to an improvement in the economic value of forecastable excess returns.

Our results have interesting implications for asset managers. We find that, despite the three ratios prove to have a weak forecasting power in a *statistical* perspective (compared to a simple average benchmark), they can still be used to select portfolio allocations that deliver better realized risk-adjusted performances than strategies that disregard predictability. In particular, our analysis indicates that MS forecast models turn out to be a useful tool to a *long-term*, *highly risk-averse* investor. Indeed, MS models are generally superior at predicting the shape of the long-run, unconditional density of excess returns, while they are in general mediocre tools to forecast the short-run dynamics. As for the performance ranking of the three alternative predictors, the earnings-price ratio turns out to be the most effective one, outperforming the popular dividend-price ratio in the majority of the exercises.

Finally, we call the attention on the fact that, when transaction costs are taken into account, they may erode any superior performance achieved thanks to predictability. However, in this paper we only consider transaction costs applied to switching strategies, where it is simple to account for them "ex-ante", such that the investor will refrain to implement switches in allocation that would yield more cost than value. Such ex-ante computations would not be trivial when more sophisticated optimization techniques were adopted to select the portfolio allocation and may represent an interesting extension to our paper.

In Famy (2007) standardised variables (i.e., that represent information about the current level of a variable relative to its recent history) appear to have more forecasting power than raw variables. It would be interesting to check the effects of standardizing our variables on the economic value of predictability. Although with reference to US data, McMillan (2011) shows that the bond-equity yield ratio has predictive power for stock returns. McMillan et al. (2014) use a range of macroeconomic variables to show that most low frequency predictability, especially at long horizons, derives from time-varying risk premia. While valuation ratios may better detect the typical decline in the equity risk premium near business cycle peaks, macroeconomic variables more readily pick up the rise in the premium later in recessions near cyclical troughs". It follows that combining both sets of predictor variables may significantly improve the overall forecasting performance.

Weigand and Irons (2007) have extensively argued in favor of the predictive power of the (inverse of the) price-earning ratio. However, they emphasize the importance of comparing short- and long-term earning-price ratios. Almadi et al. (2014) have recently proposed using diffusion indices (i.e., principal components) to summarize the complex information in a myriad of macroeconomic and valuation ratios to strengthen predictability. It would be interesting to check whether using these additional techniques to distil predictors may affect our conclusions when combined with Markov switches.

References

Almadi, H., Rapach, D. E., Suri, A. (2014). Return predictability and dynamic asset allocation: how often should investors rebalance? *Journal of Portfolio Management*, 40, 16-26.

Ang, A., Bekaert, G. (2004). How regimes affect asset allocation. *Financial Analysts Journal*, 60, 86-99.

Ang, A., Timmermann, A. (2011). Regime changes and financial markets. NBER Working paper No. 17182.

Barberis, N. (2000). Investing for the long run when returns are predictable. *Journal of Finance,* 55, 225-264.

Bekaert, G., Erb, C. B., Harvey, C. R., Viskanta, T. E. (1998). Distributional characteristics of emerging market returns and asset allocation. *Journal of Portfolio Management*, 24, 102-116.

Bossaerts, P. and Hillion, P. (1999). Implementing statistical criteria to select return forecasting models: what do we learn? *Review of Financial Studies*, 12, 405–428.

Brooks, C. (2014). *Introductory Econometrics for Finance*, Cambridge University Press 3rd edition.

Bulla, J., Mergner, S., Bulla, I., Sesboüé, A., Chesneau, C. (2011). Markov-switching asset allocation: do profitable strategies exist? *Journal of Asset Management*, 12, 310-321.

Campbell, J.Y., Shiller, R.J. (1998). Valuation ratios and the long-run stock market outlook. *Journal of Portfolio Management*, 24, 11-26.

Campbell, J.Y., Thompson, S.B. (2008). Predicting excess stock returns out of sample: can anything beat the historical average? *Review of Financial Studies*, 21, 1509-1531.

Cardinale, M., Navone, M., Pioch, A. (2014). The power of dynamic asset allocation. *Journal of Portfolio Management*, 40, 47-60.

Cochrane, J.H. (2008). The dog that did not bark: a defense of return predictability. *Review of Financial Studies*, 21, 1533-1575.

Dangl, T., Halling, M. (2012). Predictive regressions with time-varying coefficients. *Journal of Financial Economics*, 106, 157-181.

Das, S. R., & Uppal, R. (2004). Systemic risk and international portfolio choice. *Journal of Finance*, 59, 2809-2834.

Davies, R. (1977). Hypothesis testing when a nuisance parameter is present only under the alternative. *Biometrika*, 64, 247–254.

Fabozzi, F., K., Focardi, S., M., Kolm, P., N. (2006). *Financial Modeling of the Equity Market: From CAPM to Cointegration*, Wiley and Sons.

Fama, E.F., French, K.F. (1988). Dividend yields and expected stock returns. *Journal of Financial Economics*, 22, 3-25.

Famy, G. (2007). Managing market risk with conditioning information. *Journal of Asset Management*, 7, 412-418.

Fugazza, C., Guidolin, M., Nicodano, G. (2008). Diversifying in public real estate: The ex-post performance. *Journal of Asset Management*, 8, 361-373.

Gonzalo, J., Pitarakis, J. Y. (2012). Regime-specific predictability in predictive regressions. *Journal of Business and Economic Statistics*, 30, 229-241.

Goyal, A., Welch, I. (2003). Predicting the equity premium with dividend ratios. *Management Science*, 49, 639-654.

Guidolin, M. (2011a). Markov switching models in empirical finance. *Advances in econometrics* 27B, 1-86.

Guidolin, M. (2011b). Markov switching in portfolio choice and asset pricing models: A survey. *Advances in Econometrics*, 27B, 87-178.

Guidolin, M., Hyde, S. (2012). Can VAR models capture regime shifts in asset returns? A long-horizon strategic asset allocation perspective. *Journal of Banking and Finance*, 36, 695-716.

Guidolin, M., Hyde, S., McMillan, Ono, S. (2009). Non-linear predictability in stock and bond returns: when and where is it exploitable? *International Journal of Forecasting*, 25, 373-399.

Guidolin, M., Ossola, E. (2009). Do jumps matter in emerging market portfolio strategies? in *Financial Innovations in Emerging Markets* (G. Gregoriou, ed.), Chapman Hall, London, pp. 148-184.

Guidolin, M., Ria, F. (2011). Regime shifts in mean-variance efficient frontiers: some international evidence. *Journal of Asset Management*, 12, 322-349.

Guidolin, M., Timmermann, A. (2006). An econometric model of nonlinear dynamics in the joint distribution of stock and bond returns. *Journal of Applied Econometrics*, 21, 1-22.

Guidolin, M., Timmermann, A. (2007). Asset allocation under multivariate regime switching. *Journal of Economic Dynamics & Control,* 31, 3503-3544.

Hamilton, J.D. (1989). A new approach to the economic analysis of nonstationary time series and the business cycle. *Econometrica*, 57, 357-384.

Henkel, S.J., Martin, J.S., Nardari, F. (2011). Time-varying short-horizon predictability. *Journal of Financial Economics*, 99, 560-580.

Hübner, G. (2007). How do performance measures perform?. *Journal of Portfolio Management*, 33, 64-74.

Jang, B.-G., Koo, H. K., Liu, H., Loewenstein, M. (2007). Liquidity premia and transaction costs. *Journal of Finance*, 62, 2329-2366.

Jones, S. K., and Stine, J. B. (2010). Expected utility and the non-normal returns of common portfolio rebalancing strategies. *Journal of Asset Management*, 10, 406-419.

Keim, D. B., Madhavan, A. (1997). Transactions costs and investment style: an inter-exchange analysis of institutional equity trades. *Journal of Financial Economics*, 46, 265-292.

Kritzman, M., Page, S., Turkington, D. (2012). Regime shifts: implications for dynamic strategies. *Financial Analysts Journal*, 68, 22-39.

Leitch, G., Tanner, E., J. (1991). Economic forecast evaluation: profits versus the conventional error measures, *American Economic Review*, 81, 580-590.

Lesmond, D. A., Ogden, J. P., Trzcinka, C. A. (1999). A new estimate of transaction costs. *Review of Financial Studies*, 12, 1113-1141.

Lettau, M., Van Nieuwerburgh, S. (2008). Reconciling the return predictability evidence. *Review of Financial Studies*, 21, 1607-1652.

Leung, M. , Daouk, H., Chen, A. (2000). Forecasting stock indices: a comparison of classification and level estimation models, *International Journal of Forecasting*, 16, 173–190.

Lynch, A. W., Balduzzi, P. (2000). Predictability and transaction costs: the impact on rebalancing rules and behavior. *Journal of Finance*, 55, 2285-2309.

Malevergne, Y., Sornette, D. (2005). Higher-moment portfolio theory. *Journal of Portfolio Management*, 31, 49-55.

McMillan, D. (2011). Does the BEYR help predict UK sector returns?. *Journal of Asset Management*, 12, 146-156.

McMillan, D., Black, A., Klinkowska, O., McMillan, F. (2015). The existence and source of stock return predictability: evidence from dividend, output and consumption ratios. *Journal of Asset Management*, 16, 186-208.

Merton, R. C. (1969). Lifetime portfolio selection under uncertainty: the continuous-time case. *Review of Economics and Statistics*, 51, 247-257.

Nelson, C. R., Kim, M. J. (1993). Predictable stock returns: the role of small sample bias. *Journal of Finance*, 48, 641-661.

Paye, B.F., Timmermann, A. (2006). Instability of return prediction models. *Journal of Empirical Finance*, 13, 274-315.

Pesaran, M.H., Timmermann, A. (1995). Predictability of stock returns: robustness and economic significance. *Journal of Finance*, 50, 1201-1228.

Samuelson, P. A. (1969). Lifetime portfolio selection by dynamic stochastic programming. *Review of Economics and Sstatistics*, 51, 239-246.

Stambaugh, R. F. (1999). Predictive regressions. *Journal of Financial Economics*, 54, 375-421.

Stoll, H. R., Whaley, R. E. (1983). Transaction costs and the small firm effect. *Journal of Financial Economics*, 12, 57-79.

Timmermann, A. (2008). Elusive return predictability. *International Journal of Forecasting*, 24, 1-18.

Zhou, G., Rapach, D. (2013). Forecasting stock returns. *Handbook of Economic Forecasting* (Elliott, G., and Timmermann, A., eds.), Set 2A-2B. Elsevier.

Weigand, R. A., Irons, R. (2007). The market p/e ratio, earnings trends, and stock return forecasts. *Journal of Portfolio Management*, 33, 87.

Welch, I., Goyal, A. (2008). A comprehensive look at the empirical performance of equity premium prediction. *Review of Financial Studies*, 21,1455-1508.

Table 1: Summary Statistics

The table shows summary statistics for nominal stock returns, the risk-free rate, excess stock returns, and the three predictive valuation ratios analyzed. The Jarque-Bera's statistic refers to a test used to assess whether a series has a Gaussian distribution.

	Mean	St. Deviation	Median	Skewness	Kurtosis	Minimum	Maximum	Jarque-Bera
Stock Returns	0.008	0.055	0.012	-0.459	7.768	-0.348	0.357	1025.550
Risk-free Rate	0.003	0.003	0.003	1.037	1.271	0.000	0.013	317.168
Excess Returns	0.005	0.055	0.009	-0.408	7.707	-0.349	0.357	992.674
Dividend-Price Ratio	-3.338	0.455	-3.310	-0.369	-0.081	-4.524	-1.873	436.671
Earnings-Price Ratio	-2.715	0.421	-2.759	-0.729	2.788	-4.836	-1.775	94.450
Book-to-Market Ratio	0.582	0.265	0.558	0.726	1.528	0.121	2.028	185.950

Table 2: Model Selection at Different Forecast Horizons

The table shows statistics to support model selection across linear (K = 1) and non-linear Markov switching (K = 2 and 3) models.

Predictor	Horizon	N. of parameters			N.of Obs.	Saturation Ratio			Log-likelihood			Akaike IC			Hannan-Quinn IC			Bayes-Schwarz IC		
		K = 1	K = 2	K = 3		K = 1	K = 2	K = 3	K = 1	K = 2	K = 3	K = 1	K = 2	K = 3	K = 1	K = 2	K = 3	K = 1	K = 2	K = 3
Dividend-price ratio	1 month	3	8	15	1043	347.7	130.4	69.5	1546.54	1729.77	1741.59	-2.9598	-3.302	**-3.311**	-2.9544	**-3.287**	-3.284	-2.9456	**-3.264**	-3.240
	6 months	3	8	15	1038	346.0	129.8	69.2	1538.58	1721.09	1742.03	-2.9587	-3.301	**-3.328**	-2.9533	-3.286	**-3.301**	-2.9444	**-3.263**	-3.256
	5 years	3	8	15	984	328.0	123.0	65.6	1468.66	1642.48	1662.79	-2.979	-3.322	**-3.349**	-2.9733	-3.307	**-3.321**	-2.9641	**-3.282**	-3.275
Eearning-price ratio	1 month	3	8	15	1043	347.7	130.4	69.5	1547.21	1725.57	1747.90	-2.9611	-3.294	**-3.323**	-2.956	-3.279	**-3.296**	-2.947	**-3.256**	-3.252
	6 months	3	8	15	1038	346.0	129.8	69.2	1538.34	1716.91	1730.75	-2.9583	-3.292	**-3.306**	-2.953	-3.278	**-3.279**	-2.944	**-3.255**	-3.234
	5 years	3	8	15	984	328.0	123.0	65.6	1468.29	1641.59	1657.04	-2.9782	-3.320	**-3.338**	-2.973	-3.305	**-3.309**	-2.9633	**-3.281**	-3.263
Book-to-mkt ratio	1 month	3	8	15	1043	347.7	130.4	69.5	1547.69	1728.36	1739.90	-2.962	-3.299	**-3.308**	-2.9566	**-3.285**	-3.281	-2.9478	**-3.261**	-3.236
	6 months	3	8	15	1038	346.0	129.8	69.2	1539.79	1718.63	1739.69	-2.9611	-3.296	**-3.323**	-2.9556	-3.282	**-3.296**	-2.9468	**-3.258**	-3.252
	5 years	3	8	15	984	328.0	123.0	65.6	1468.07	1640.62	1655.77	-2.9778	-3.318	**-3.335**	-2.9721	-3.303	**-3.307**	-2.9629	**-3.279**	-3.260

Table 3: Linear and Markov Switching Model Estimates at Different Forecast Horizons

The table shows coefficient estimates for linear and Markov switching models with two states:

$$R_{t+1} = \alpha_{S_{t+1}} + \beta_{S_{t+1}} X_t + \varepsilon_{t+1} \varepsilon_{t+1} \sim (0, \sigma^2_{S_{t+1}}).$$

Horizon		α			β			Std.error			Ergodic Probabilities		Duration (months)		Transition Matrix		
		Linear	Markov switching		Linear	Markov switching		Linear	Markov switching								
			Regime 1	Regime 2		Regime 1	Regime 2		Regime 1	Regime 2	Regime 1	Regime 2	Regime 1	Regime 2		**Regime 1**	**Regime 2**
Dividend-Price Ratio																	
	Estimate	**0.0275**	**0.1515**	**0.0348**	**0.0068**	**0.0572**	**0.0079**								**Regime 1**	**0.922**	0.078
1 month	Std. Err	0.0126	0.0915	0.0103	0.0037	0.0310	0.0030	0.0550	0.1150	0.0390	0.1135	0.8865	12.87	100.53	**Regime 2**	0.010	**0.990**
	t-stat	2.1788	1.6560	3.3875	1.8146	1.8447	2.6152										
	Estimate	**-0.0309**	0.0931	0.0386	**0.0146**	0.0365	**0.0091**								**Regime 1**	**0.927**	0.073
6 month	Std. Err	0.0159	0.0784	0.0102	0.0035	0.0258	0.0030	0.0473	0.1159	0.0389	0.1162	0.8838	13.63	103.63	**Regime 2**	0.010	**0.990**
	t-stat	-1.9434	1.1878	0.0030	4.2315	1.4166	2.9963										
	Estimate	0.0182	**0.0320**	-0.0428	0.0042	-0.0084	**0.0074**								**Regime 1**	**0.913**	0.087
5 years	Std. Err	0.0155	0.0103	0.0889	0.0039	0.0278	0.0031	0.0539	0.1218	0.0388	0.0953	0.9047	11.49	109.08	**Regime 2**	0.009	**0.991**
	t-stat	1.1742	3.0965	-0.4812	1.0681	-0.3017	2.4051										
Earning-Price Ratio																	
	Estimate	**0.0284**	0.0466	**0.0193**	**0.0087**	0.0231	0.0042								**Regime 1**	**0.920**	0.080
1 month	Std. Err	0.0111	0.0851	0.0103	0.0040	0.0304	0.0038	0.0549	0.1158	0.0389	0.1178	0.8822	12.52	93.74	**Regime 2**	0.011	**0.989**
	t-stat	2.5569	0.5472	1.8742	2.1500	0.7611	1.0820										
	Estimate	0.0221	0.0700	**0.0203**	**0.0080**	0.0315	0.0045								**Regime 1**	**0.923**	0.077
6 month	Std. Err	0.0156	0.1033	0.0088	0.0042	0.0369	0.0033	0.0539	0.1163	0.0389	0.1178	0.8822	13.04	97.67	**Regime 2**	0.010	**0.990**
	t-stat	1.4167	0.6770	2.3178	1.9174	0.8542	1.3978										
	Estimate	0.0125	**0.0219**	-0.1642	0.0030	-0.0534	0.0054								**Regime 1**	**0.912**	0.088
5 years	Std. Err	0.0152	0.0093	0.1169	0.0046	0.0419	0.0035	0.0539	0.1211	0.0389	0.0947	0.9053	11.38	108.79	**Regime 2**	0.009	**0.991**
	t-stat	0.8224	2.3513	-1.4040	0.6564	-1.2725	1.5752										
Book-to-Market Ratio																	
	Estimate	-0.0040	**-0.0773**	**0.0067**	**0.0152**	**0.0799**	0.0029								**Regime 1**	**0.911**	0.089
1 month	Std. Err	0.0041	0.0254	0.0036	0.0064	0.0296	0.0057	0.0549	0.1118	0.0387	0.1190	0.8810	11.19	82.86	**Regime 2**	0.012	**0.988**
	t-stat	-0.9813	-3.0442	1.8901	2.3623	2.7012	0.5084										
	Estimate	**-0.0904**	**-0.0578**	0.0042	**0.0223**	**0.0588**	0.0070								**Regime 1**	**0.920**	0.080
6 month	Std. Err	0.0105	0.0220	0.0036	0.0064	0.0277	0.0057	0.0473	0.1134	0.0388	0.1207	0.8793	12.43	90.60	**Regime 2**	0.011	**0.989**
	t-stat	-8.6095	-2.6307	1.1636	3.5000	2.1198	1.2224										
	Estimate	0.0039	-0.0050	0.0027	0.0009	-0.0183	0.0078								**Regime 1**	**0.912**	0.088
5 years	Std. Err	0.0096	0.0235	0.0038	0.0067	0.0317	0.0059	0.0539	0.1219	0.0389	0.0946	0.9054	11.36	108.75	**Regime 2**	0.009	**0.991**
	t-stat	0.4063	-0.2114	0.7034	0.1285	-0.5766	1.3330										

Note: boldfaced coefficients are significant with a p-value of 0.1 or less.

Table 4: Switching Strategies Realized Out-of-Sample Performance Indicators

The table shows realized performances from a *pseudo*, recursive out-of-sample exercise at different levels of transaction costs, in which switching strategies are alternatively based on linear vs. non-linear (Markov switching) predictive regressions. In the table, B&H means "buy-and-hold".

	Horizon	1 month			6 months			60 months		
		MS	**Linear**	**B&H**	**MS**	**Linear**	**B&H**	**MS**	**Linear**	**B&H**
Dividend-Price Ratio					**Hit Ratios**					
		0.4950	0.4130	__	0.5140	0.4980	__	0.6750	0.5990	__
					No Transaction Costs					
	Mean	0.0070	0.0073	0.0076	0.0069	0.0073	0.0076	0.0069	0.0066	0.0077
	Variance	0.0012	0.0013	0.0018	0.0011	0.0013	0.0018	0.0012	0.0010	0.0018
	St. Deviation	0.0351	0.0365	0.0427	0.0336	0.0362	0.0427	0.0352	0.0316	0.0429
	Sharpe Ratio	0.0826	0.0880	0.0837	0.0855	0.0892	0.0837	0.0822	0.0819	0.0839
					Low Transaction Costs					
	Mean	0.0050	0.0044	0.0076	0.0056	0.0070	0.0076	0.0067	0.0069	0.0077
	Variance	0.0009	0.0006	0.0018	0.0008	0.0012	0.0018	0.0009	0.0010	0.0018
	St. Deviation	0.0301	0.0246	0.0427	0.0289	0.0351	0.0427	0.0307	0.0311	0.0429
	Sharpe Ratio	0.0326	0.0143	0.0837	0.0545	0.0836	0.0837	0.0874	0.0907	0.0839
					High Transaction Costs					
	Mean	0.0037	0.0045	0.0076	0.0055	0.0063	0.0076	0.0056	0.0067	0.0077
	Variance	0.0007	0.0003	0.0018	0.0006	0.0011	0.0018	0.0007	0.0010	0.0018
	St. Deviation	0.0256	0.0183	0.0427	0.0246	0.0333	0.0427	0.0259	0.0310	0.0429
	Sharpe Ratio	-0.0158	0.0257	0.0837	0.0591	0.0661	0.0837	0.0591	0.0840	0.0839
Earning-Price Ratio					**Hit Ratios**					
		0.5290	0.5480	__	0.6010	0.5360	__	0.7140	0.5550	__
					No Transaction Costs					
	Mean	0.0084	0.0077	0.0076	0.0079	0.0077	0.0076	0.0079	0.0083	0.0077
	Variance	0.0015	0.0011	0.0018	0.0015	0.0011	0.0018	0.0017	0.0014	0.0018
	St. Deviation	0.0392	0.0334	0.0427	0.0386	0.0338	0.0427	0.0415	0.0371	0.0429
	Sharpe Ratio	0.1101	0.1088	0.0837	0.1001	0.1077	0.0837	0.0926	0.1150	0.0839
					Low Transaction Costs					
	Mean	0.0060	0.0054	0.0076	0.0074	0.0071	0.0076	0.0078	0.0082	0.0077
	Variance	0.0015	0.0008	0.0018	0.0014	0.0011	0.0018	0.0016	0.0014	0.0018
	St. Deviation	0.0386	0.0281	0.0427	0.0374	0.0331	0.0427	0.0399	0.0371	0.0429
	Sharpe Ratio	0.0516	0.0468	0.0837	0.0883	0.0920	0.0837	0.0948	0.1120	0.0839
					High Transaction Costs					
	Mean	0.0042	0.0039	0.0076	0.0070	0.0071	0.0076	0.0073	0.0082	0.0077
	Variance	0.0013	0.0006	0.0018	0.0010	0.0010	0.0018	0.0011	0.0014	0.0018
	St. Deviation	0.0361	0.0238	0.0427	0.0319	0.0321	0.0427	0.0328	0.0368	0.0429
	Sharpe Ratio	0.0026	-0.0056	0.0837	0.0906	0.0940	0.0837	0.0977	0.1126	0.0839
Book-to-Market Ratio					**Hit Ratios**					
		0.5180	0.4820	__	0.5890	0.4760	__	0.7240	0.5670	__
					No Transaction Costs					
	Mean	0.0080	0.0069	0.0076	0.0078	0.0070	0.0076	0.0076	0.0076	0.0077
	Variance	0.0016	0.0010	0.0018	0.0016	0.0011	0.0018	0.0017	0.0014	0.0018
	St. Deviation	0.0401	0.0320	0.0427	0.0402	0.0333	0.0427	0.0407	0.0370	0.0429
	Sharpe Ratio	0.0975	0.0878	0.0837	0.0942	0.0871	0.0837	0.0883	0.0945	0.0839
					Low Transaction Costs					
	Mean	0.0052	0.0046	0.0076	0.0072	0.0064	0.0076	0.0070	0.0072	0.0077
	Variance	0.0015	0.0007	0.0018	0.0016	0.0010	0.0018	0.0015	0.0013	0.0018
	St. Deviation	0.0391	0.0263	0.0427	0.0399	0.0317	0.0427	0.0390	0.0366	0.0429
	Sharpe Ratio	0.0285	0.0218	0.0837	0.0781	0.0740	0.0837	0.0757	0.0852	0.0839
					High Transaction Costs					
	Mean	0.0033	0.0037	0.0076	0.0062	0.0060	0.0076	0.0055	0.0072	0.0077
	Variance	0.0014	0.0006	0.0018	0.0014	0.0010	0.0018	0.0009	0.0013	0.0018
	St. Deviation	0.0378	0.0237	0.0427	0.0370	0.0309	0.0427	0.0307	0.0365	0.0429
	Sharpe Ratio	-0.0213	-0.0134	0.0837	0.0569	0.0615	0.0837	0.0471	0.0859	0.0839

Table 5: Mean-Variance Realized Out-of-Sample Performances

The table shows realized performances from a *pseudo*, recursive out-of-sample exercise in which mean-variance strategies are alternatively based on linear vs. non-linear (Markov switching) predictive regressions. In the table, CEV is the certainty equivalent realized value function. γ is the constant coefficient of risk aversion.

	Horizon	1 month		6 months		60 months	
		MS	**Linear**	**MS**	**Linear**	**MS**	**Linear**
Dividend-Price Ratio				γ = 2			
	Δ CEV	-0.0005	-0.0001	-0.0060	-0.0036	**-0.1153**	**-0.0405**
	Variance with DP	0.0009	0.0008	0.0012	0.0067	0.0147	0.0779
	Variance Gaussian IID	0.0018	0.0018	0.0123	0.0123	0.1126	0.1126
	Test Statistic	-0.2522	-0.0514	-1.3455	-0.6802	-8.0743	-2.3181
				γ = 5			
	Δ CEV	0.0000	0.0005	0.0015	0.0015	-0.0220	**-0.0400**
	Variance with DP	0.0007	0.0005	0.0004	0.0032	0.0143	0.0699
	Variance Gaussian IID	0.0010	0.0010	0.0061	0.0061	0.0673	0.0673
	Test Statistic	0.0000	0.3335	0.4862	0.4050	-1.9240	-2.6976
				γ = 10			
	Δ CEV	-0.0004	0.0002	0.0003	0.0008	**0.0501**	**-0.0624**
	Variance with DP	0.0003	0.0001	0.0002	0.0011	0.0142	0.0553
	Variance Gaussian IID	0.0003	0.0003	0.0017	0.0017	0.0487	0.0487
	Test Statistic	-0.4219	0.2495	0.1774	0.3937	4.9910	-4.8335
Earning-Price Ratio				γ = 2			
	Δ CEV	0.0002	0.0005	-0.0046	0.0027	**-0.1130**	-0.0061
	Variance with EP	0.0015	0.0013	0.0018	0.0075	0.0145	0.0929
	Variance Gaussian IID	0.0018	0.0018	0.0123	0.0123	0.1126	0.1126
	Test Statistic	0.0912	0.2353	-1.0084	0.4997	-7.9177	-0.3362
				γ = 5			
	Δ CEV	0.0002	0.0002	0.0022	0.0005	-0.0210	-0.0196
	Variance with EP	0.0011	0.0007	0.0004	0.0045	0.0142	0.0766
	Variance Gaussian IID	0.0010	0.0010	0.0061	0.0061	0.0673	0.0673
	Test Statistic	0.1141	0.1269	0.7089	0.1265	-1.8374	-1.2907
				γ = 10			
	Δ CEV	-0.0011	-0.0001	0.0008	-0.0004	**0.0508**	**-0.0627**
	Variance with EP	0.0006	0.0002	0.0003	0.0016	0.0141	0.0627
	Variance Gaussian IID	0.0003	0.0003	0.0017	0.0017	0.0487	0.0487
	Test Statistic	-0.9673	-0.1122	0.4716	-0.1813	5.0624	-4.6926
Book-to-Market Ratio				γ = 2			
	Δ CEV	-0.0004	**0.0387**	-0.0056	-0.0053	**-0.1028**	-0.0256
	Variance with BM	0.0016	0.0009	0.0021	0.0058	0.0183	0.0750
	Variance Gaussian IID	0.0018	0.0018	0.0123	0.0123	0.1126	0.1126
	Test Statistic	-0.1797	19.5183	-1.2154	-1.0260	-7.0976	-1.4765
				γ = 5			
	Δ CEV	-0.0009	-0.0007	0.0019	-0.0039	-0.0146	**-0.0290**
	Variance with BM	0.0013	0.0006	0.0005	0.0043	0.0146	0.0662
	Variance Gaussian IID	0.0010	0.0010	0.0061	0.0061	0.0673	0.0673
	Test Statistic	-0.4908	-0.4577	0.6098	-0.9958	-1.2740	-1.9827
				γ = 10			
	Δ CEV	-0.0014	-0.0016	0.0006	**-0.0076**	**0.0557**	**-0.0408**
	Variance with BM	0.0006	0.0004	0.0003	0.0026	0.0139	0.0532
	Variance Gaussian IID	0.0003	0.0003	0.0017	0.0017	0.0487	0.0487
	Test Statistic	-1.2229	-1.5816	0.3527	-3.0178	5.5594	-3.1928

Figure 1: Dynamics of Coefficient Estimates at 1 Month Horizon for the Dividend-Price, Earning-Price, and Book-to-Market Ratios

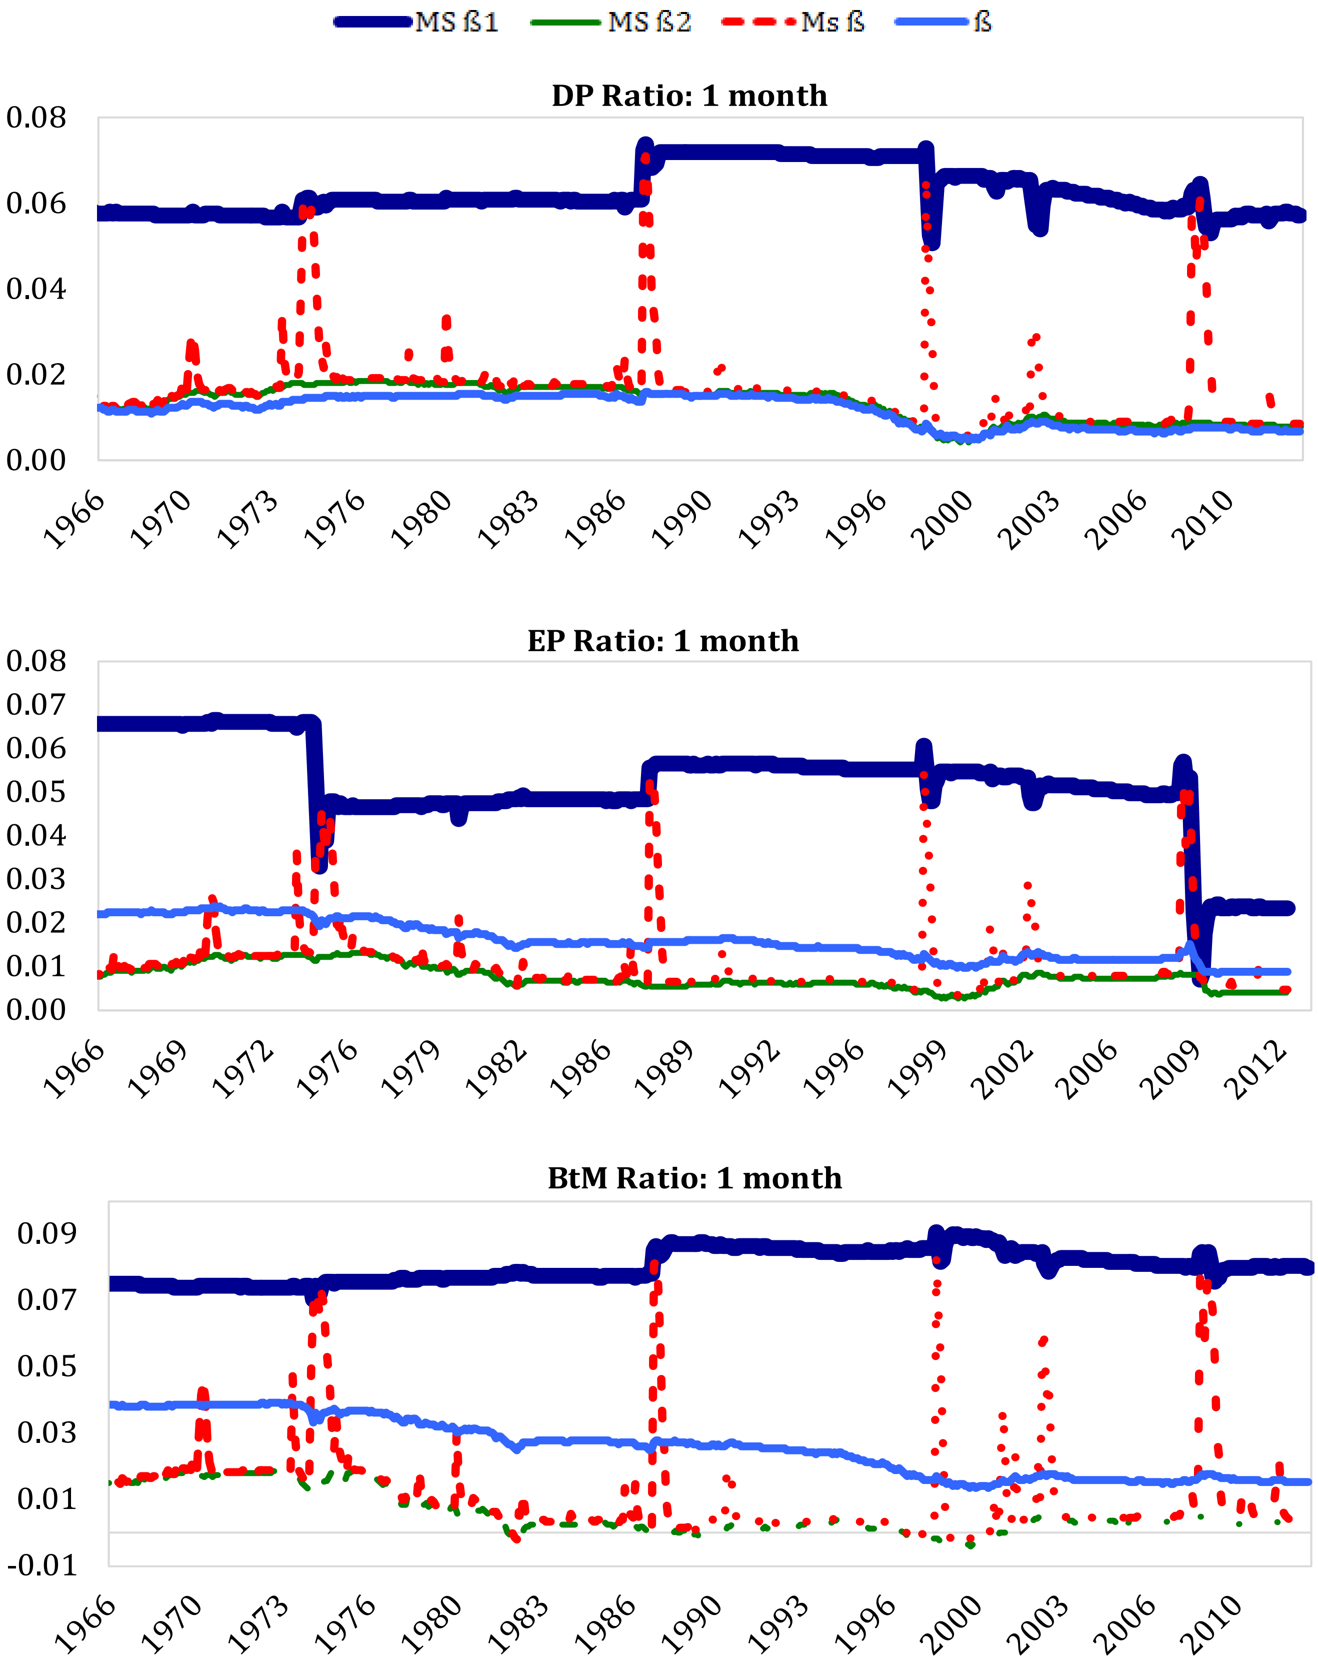

Figure 2a: Dynamics of the Differences in Root-Mean-Square Forecast Errors Implied by Dividend-Price Ratio Predictability

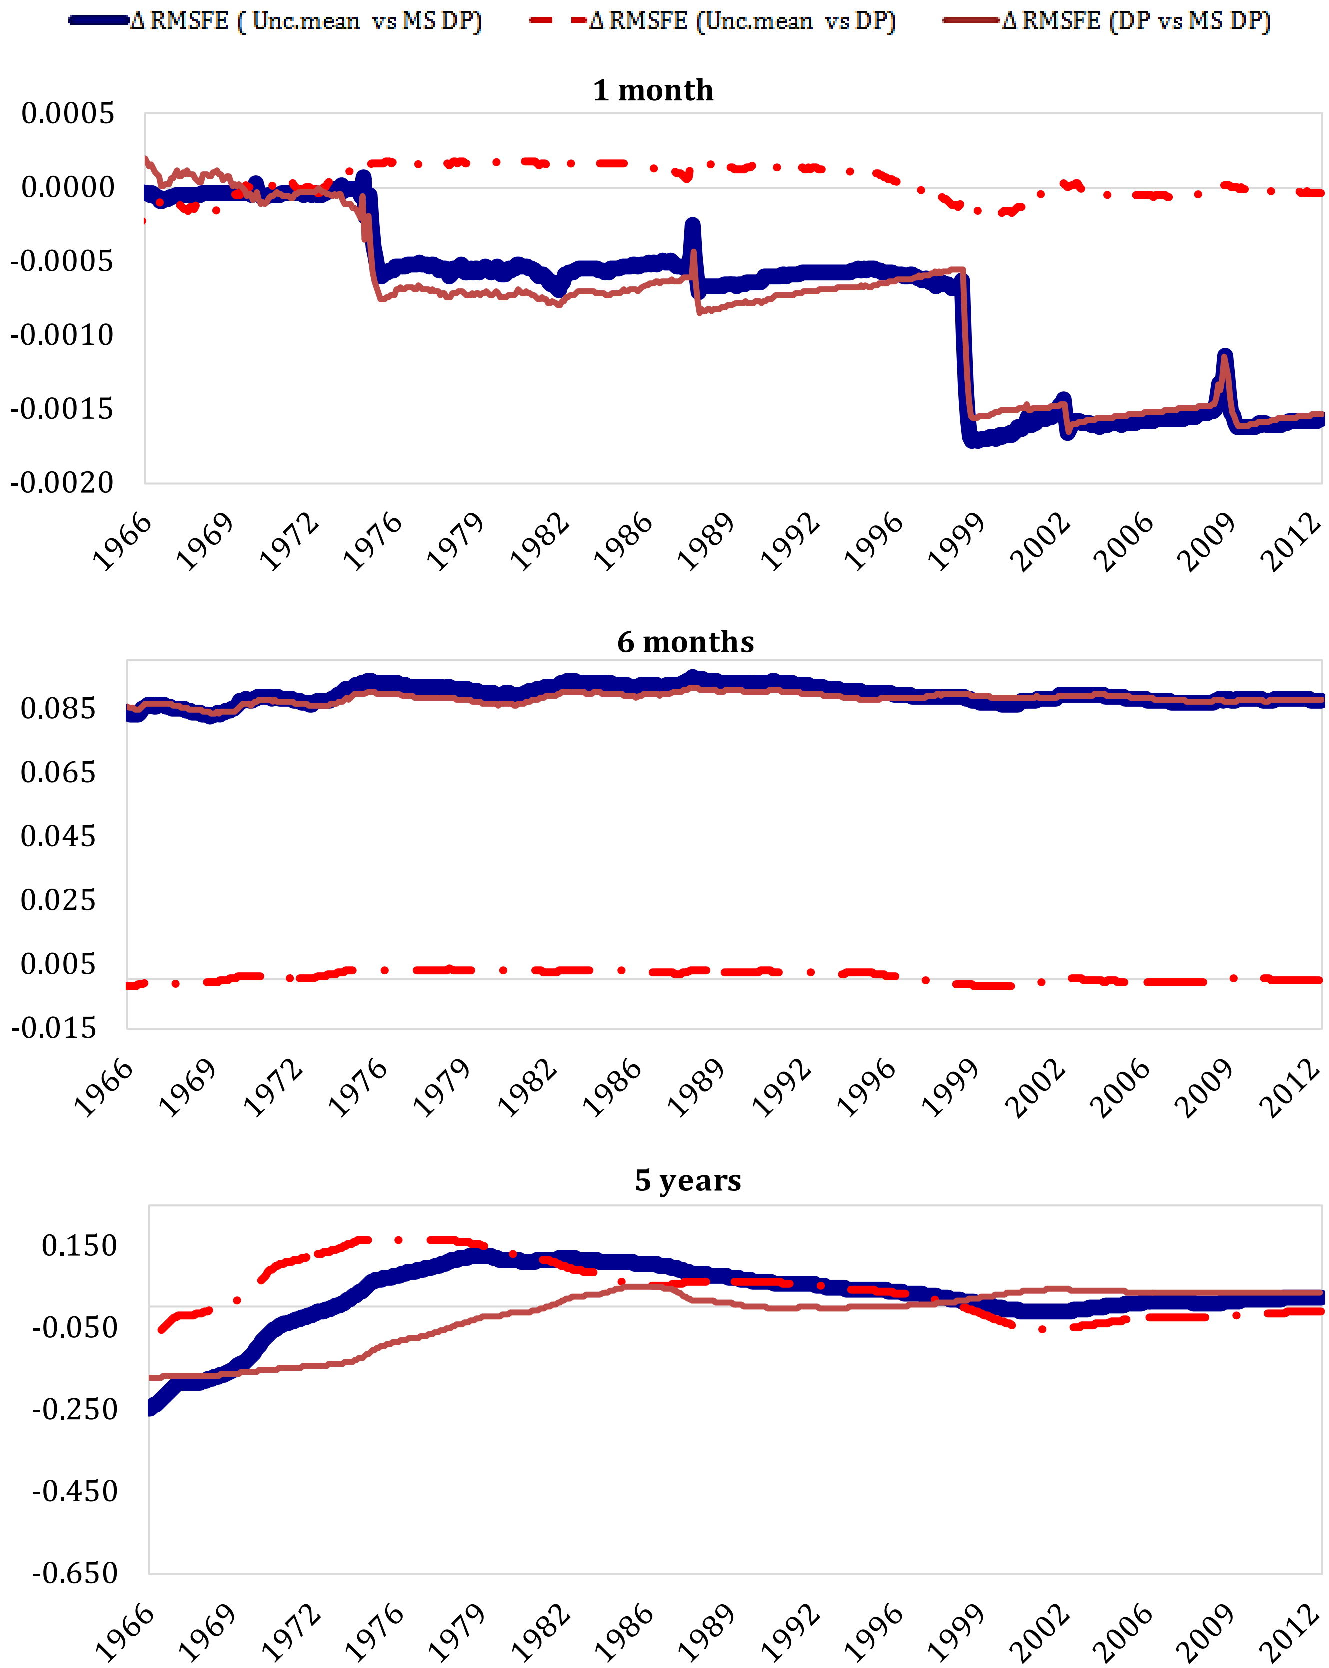

Figure 2b: Dynamics of the Differences in Root-Mean-Square Forecast Errors Implied by Earnings-Price Ratio Predictability

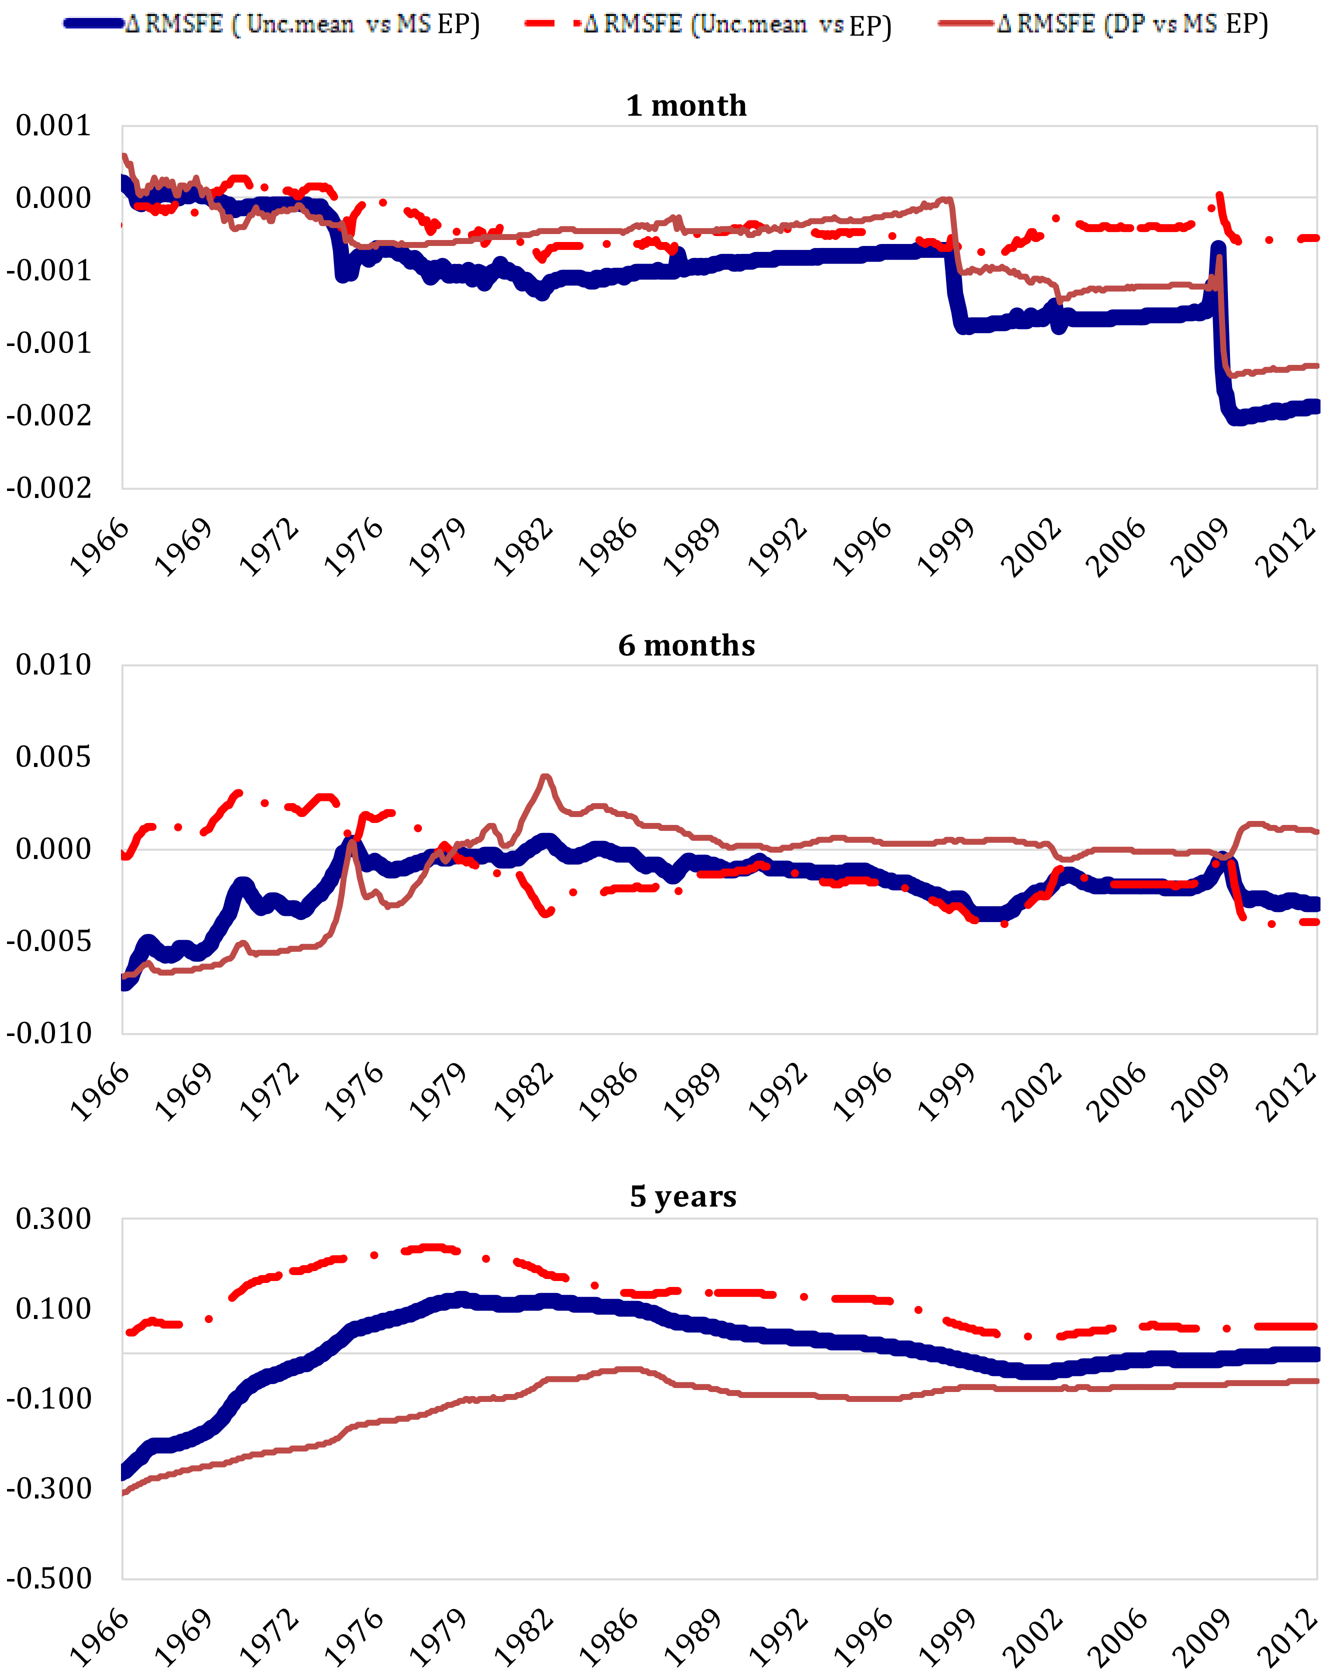

Figure 2c: Dynamics of the Differences in Root-Mean-Square Forecast Errors Implied by Book-to-Market Ratio Predictability

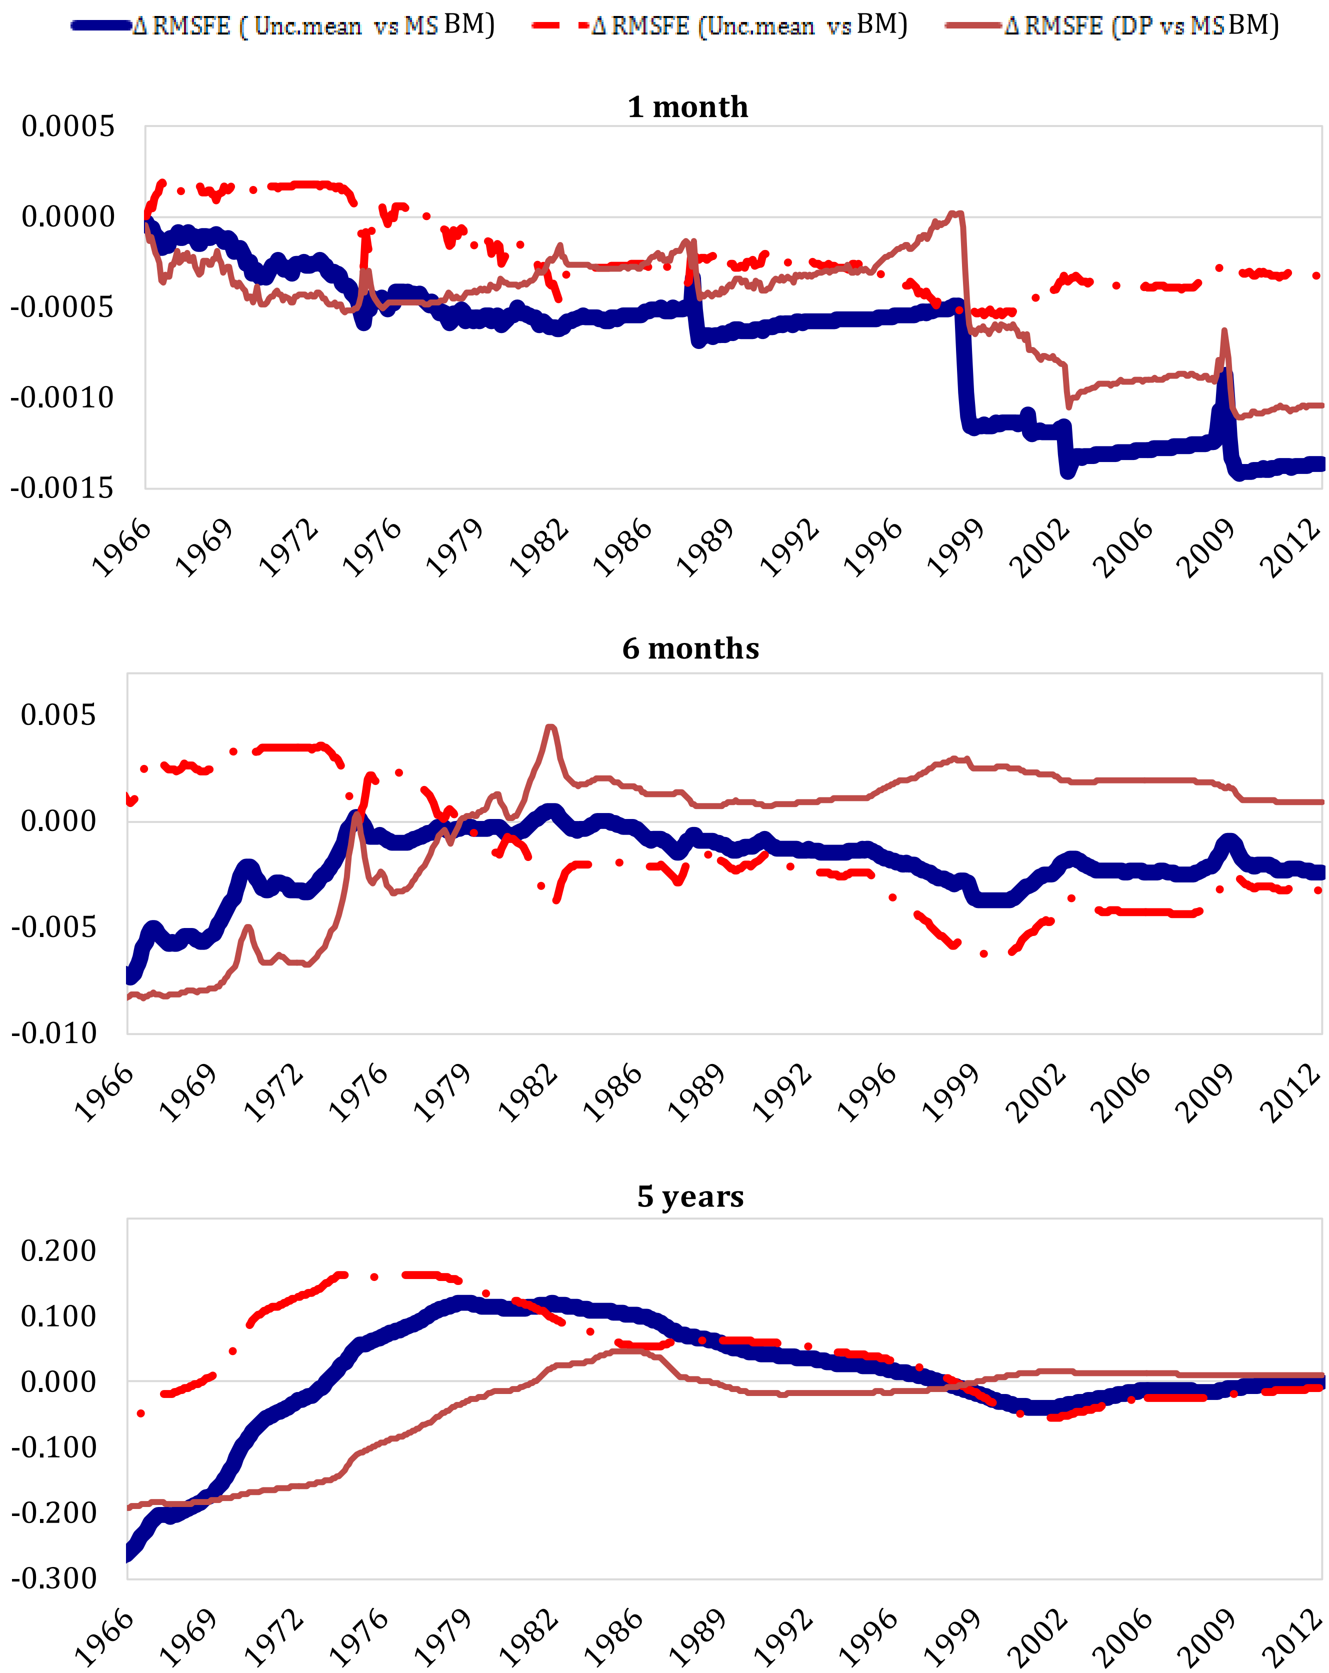

Figure 3a: Increase in Mean-Variance Certainty Equivalent Return from Exploiting Dividend-Price Ratio Predictability at Different Investment Horizons in the MS Model

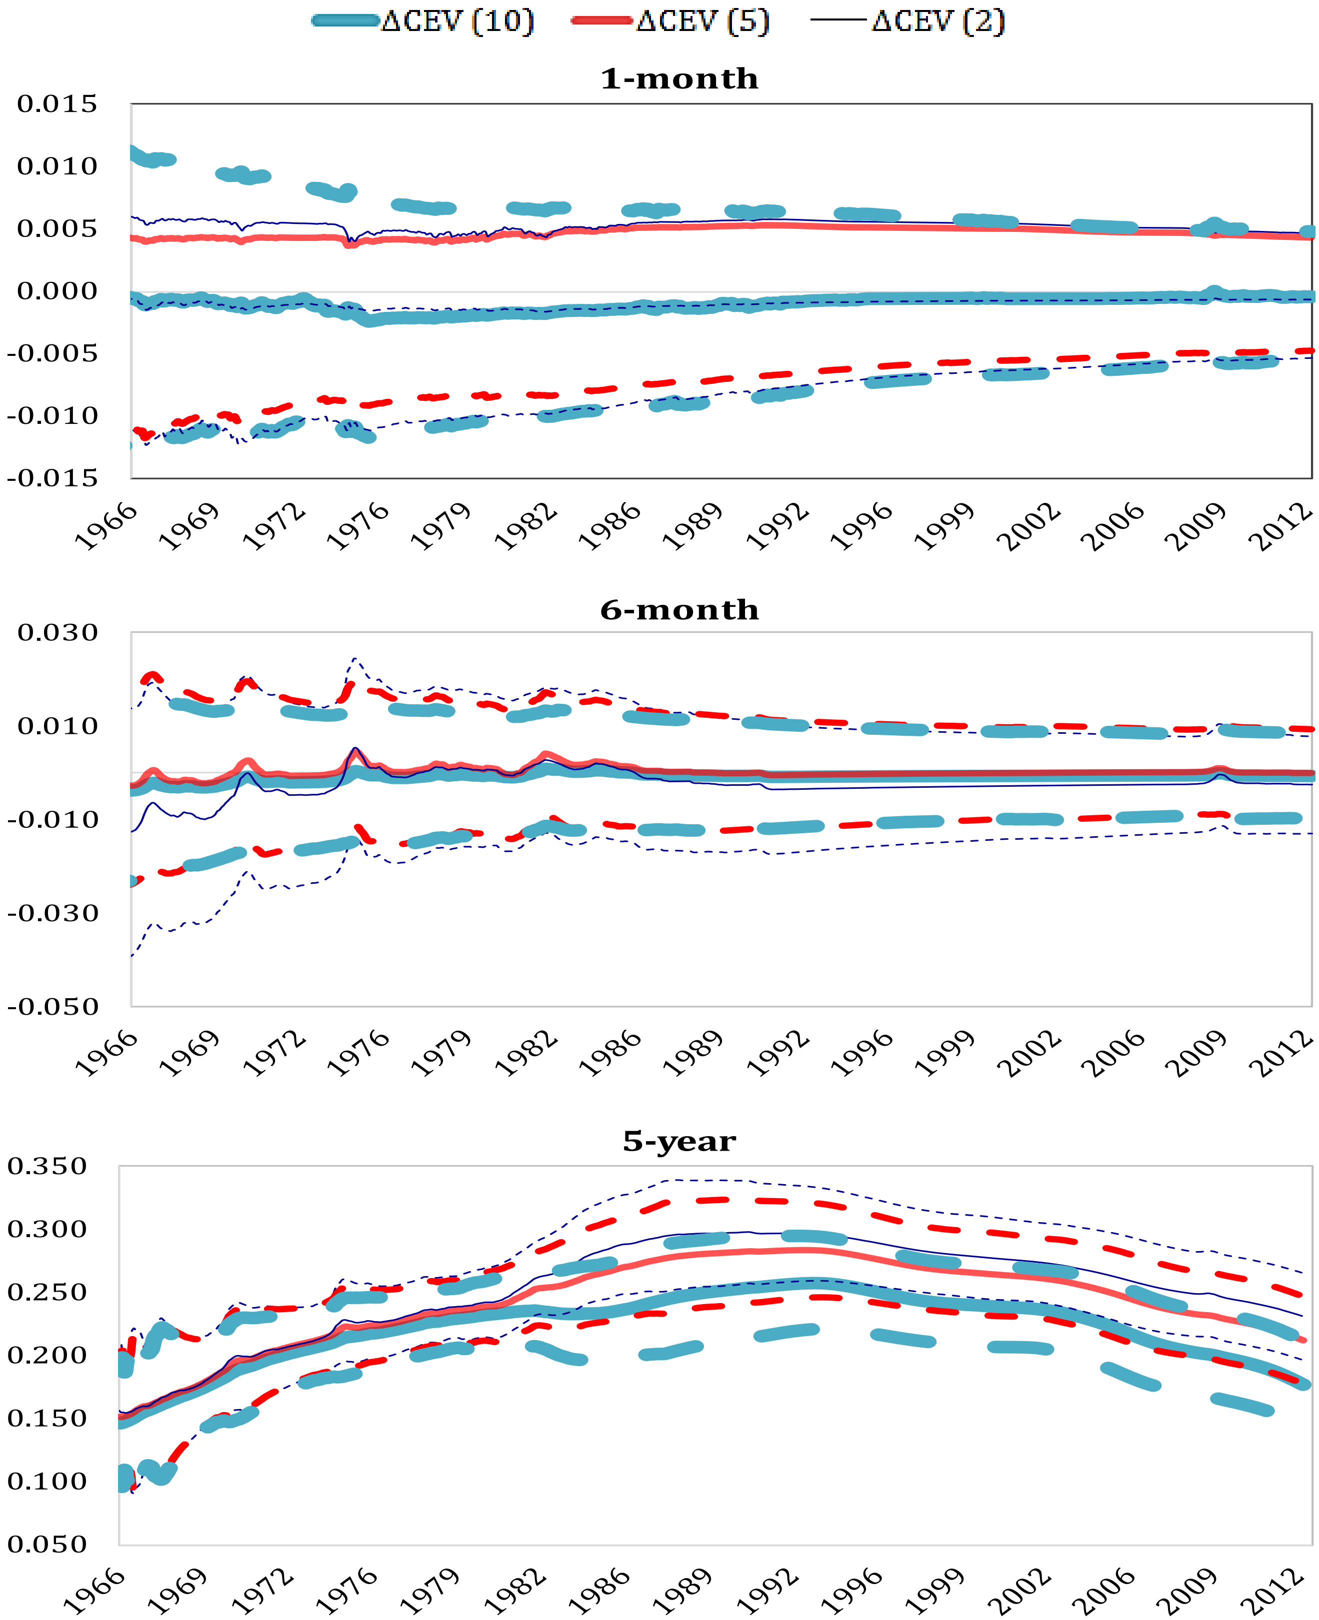

Figure 3b: Increase in Mean-Variance Certainty Equivalent Return from Exploiting Earnings Price Ratio Predictability at Different Investment Horizons in the MS Model

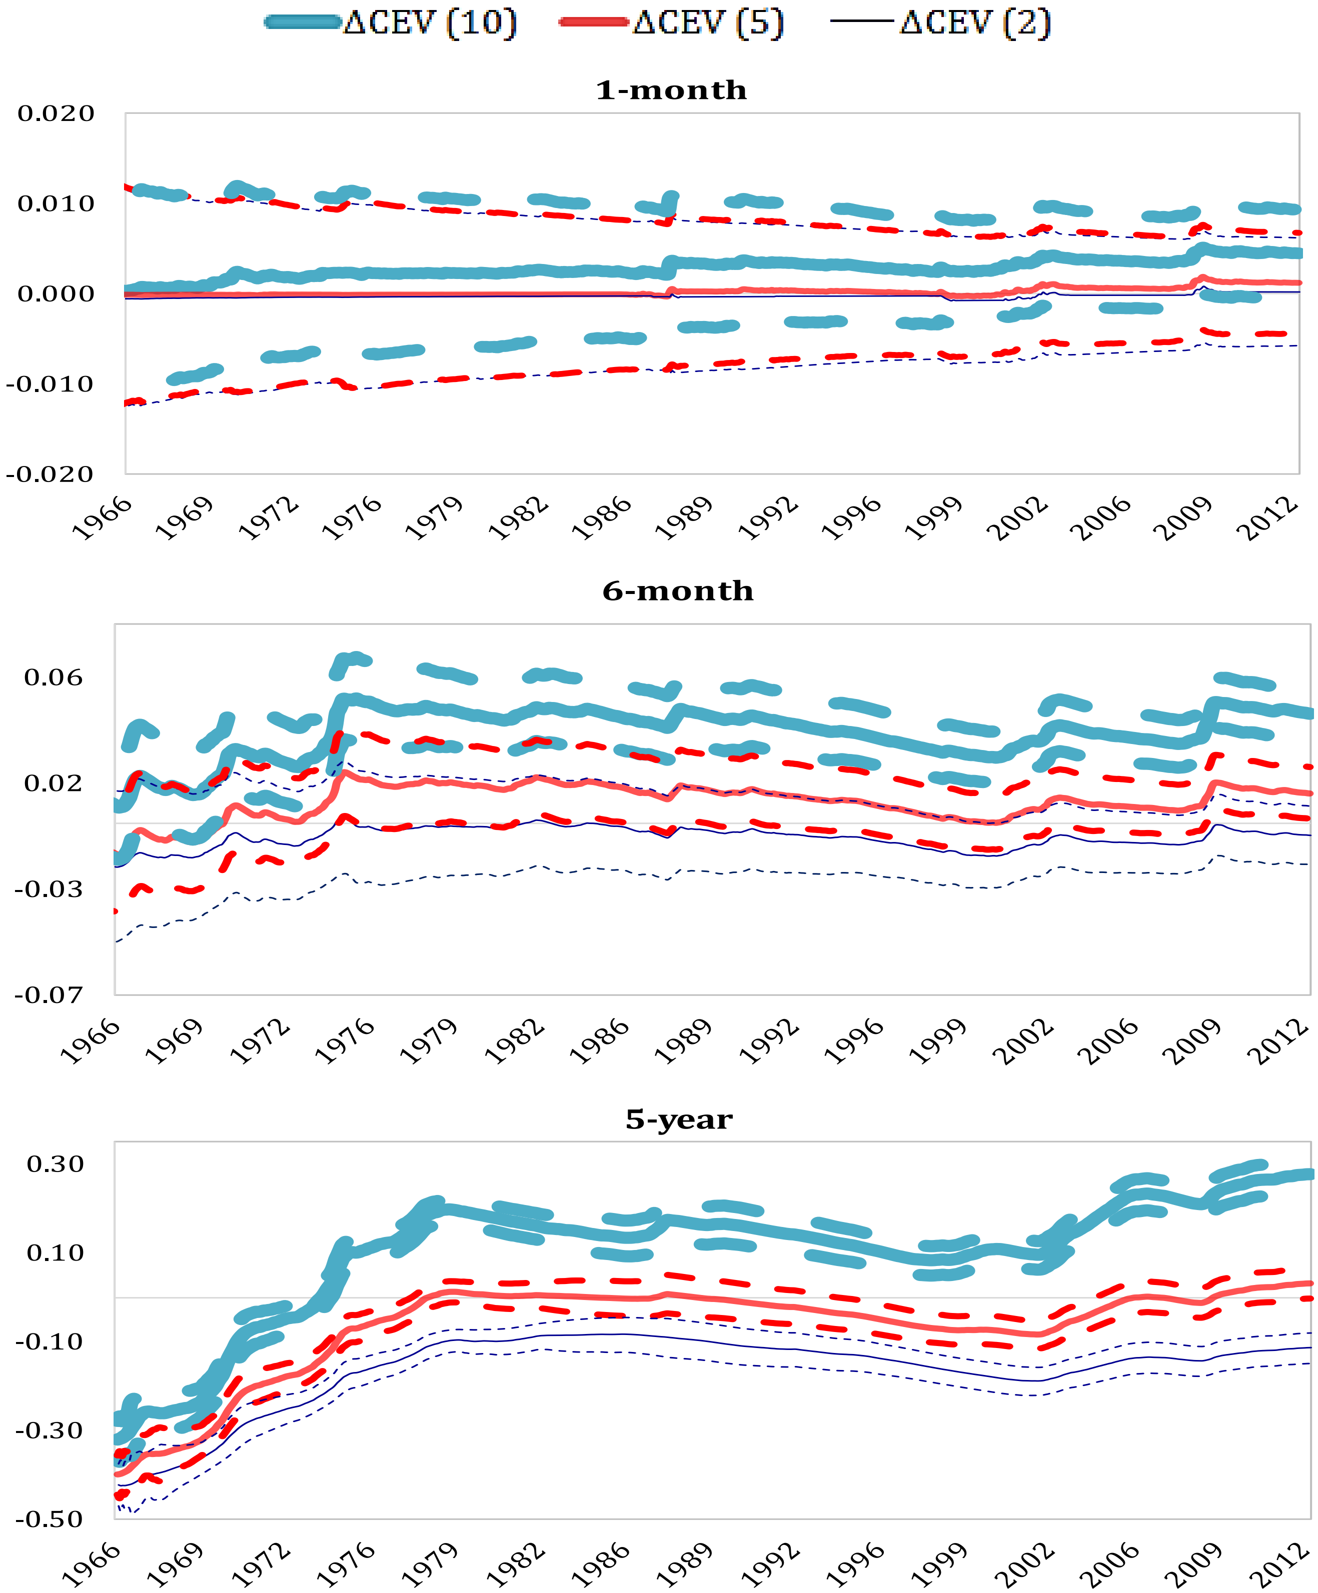

Figure 3c: Increase in Mean-Variance Certainty Equivalent Return from Exploiting Book to Market Ratio Predictability at Different Investment Horizons in the MS Model

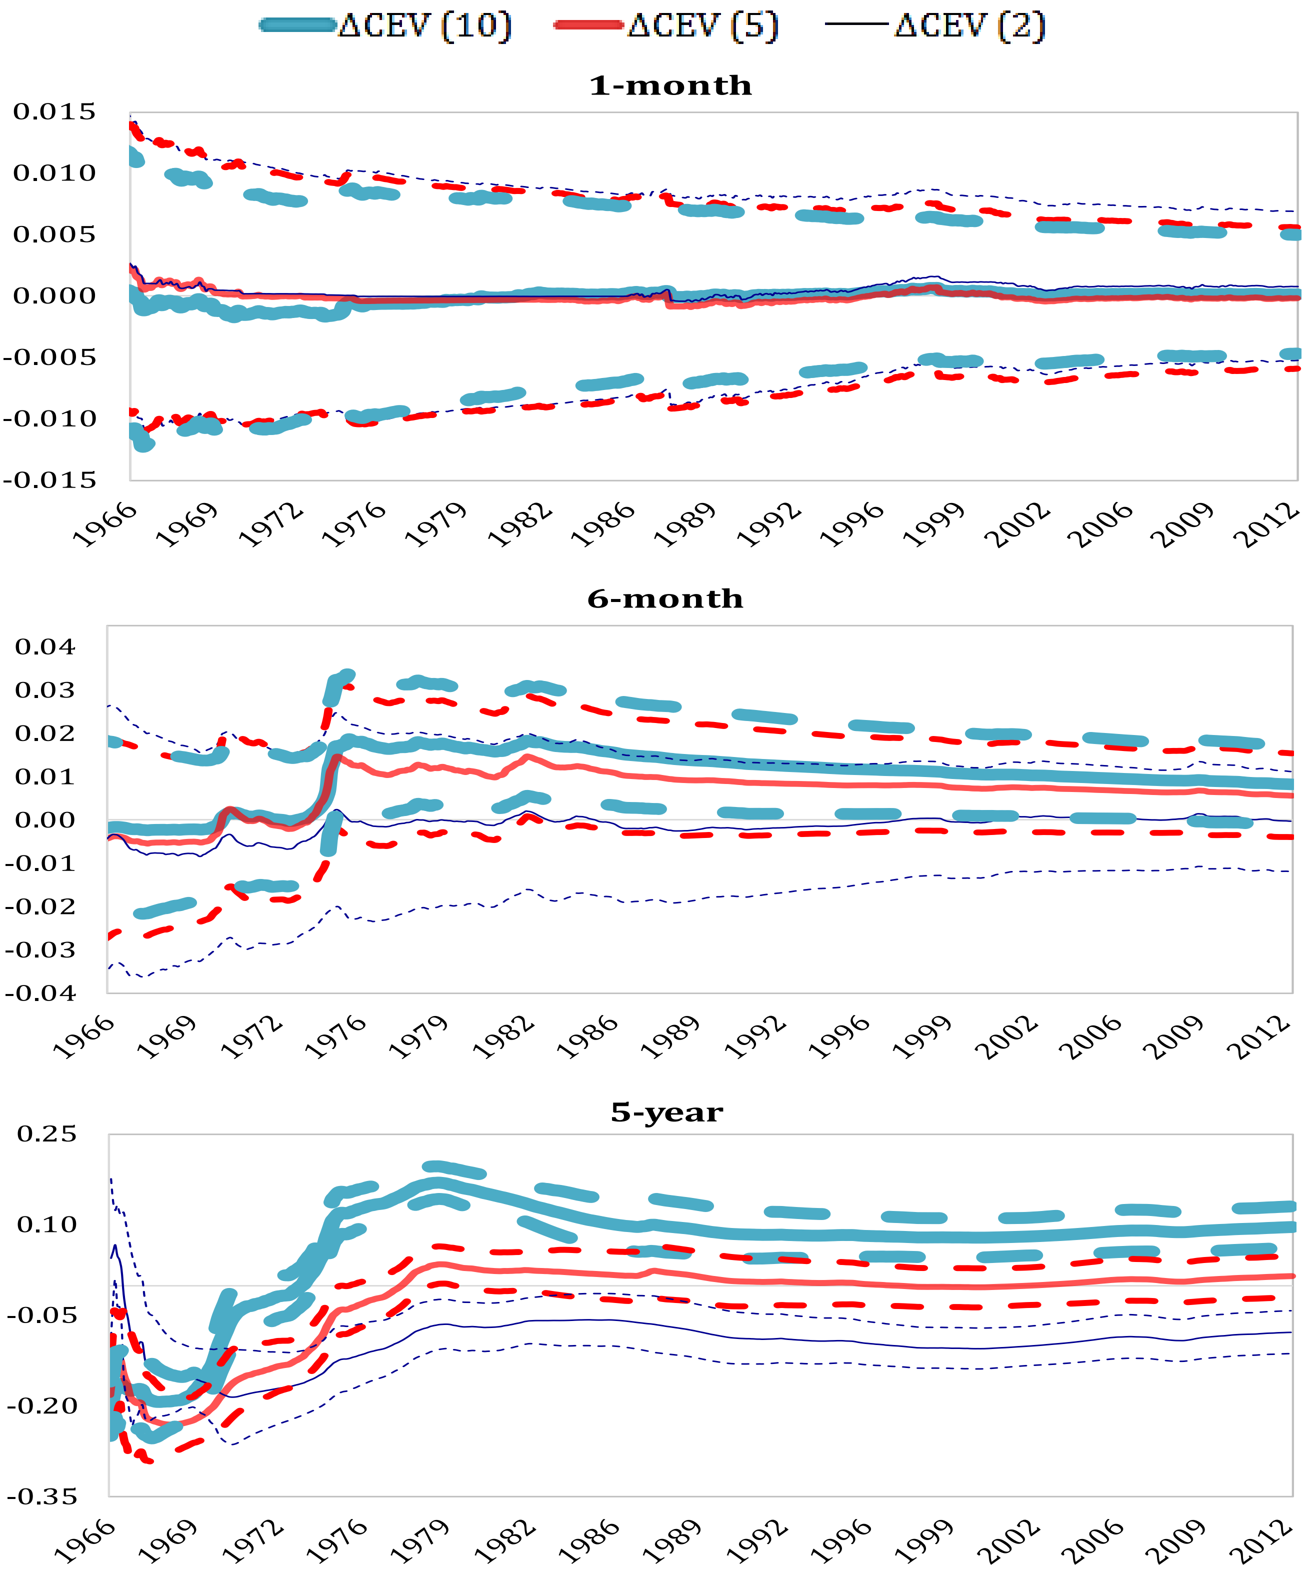

Figure 4: Annualized Compensatory Fee at Different Horizons Due to Switching from Linear to MS Models under Buy-and-Hold and Rebalancing Strategies for Power Utility, CRRA Preferences (γ = 5)