7-2017

Forecasting stock returns in good and bad times: The role of market states

Dashan HUANG
Singapore Management University, DASHANHUANG@smu.edu.sg

Fuwei JIANG
Central University of Finance and Economics

Jun TU
Singapore Management University, tujun@smu.edu.sg

Guofu ZHOU
Washington University in St. Louis

Forecasting Stock Returns in Good and Bad Times: The Role of Market States*

Dashan Huang

Singapore Management University

Fuwei Jiang

Central University of Finance and Economics

Jun Tu

Singapore Management University

Guofu Zhou

Washington University in St. Louis

Current version: July 2017

*We thank Michael Brennan, Jie (Jay) Cao, Zhi Da, Ohad Kadan, Raymond Kan, Hong Liu, Roger Loh, Ernst Maug, Massimo Massa, Pavel Savor, Wenlan Qian, Farris Shuggi, Elvira Sojli, Hitoshi Takehara, Ngoc-Khanh Tran, Joe Zhang, and seminar participants at Beijing University, London Imperial College, Shanghai Advanced Institute of Finance, Southwestern University of Finance and Economics, Tsinghua University, University of Houston, University of Texas–Dallas, University of Warwick, Washington University in St. Louis, SMU Finance Summer Camp 2014, China International Conference in Finance 2014 , World Finance & banking Symposium 2014, and 27th Australasian Finance & Banking Conference for helpful comments. We thank Shapour Mohammadi for making the quantile regression MatLab code available, Amit Goyal for making his data available, and Xinrui Duan for excellent research assistance.

Send correspondence to Guofu Zhou, Olin School of Business, Washington University in St. Louis, St. Louis, MO 63130; telephone: 314-935–6384. E-mail: zhou@wustl.edu.

Forecasting Stock Returns in Good and Bad Times: The Role of Market States

Abstract

This paper proposes a two-state predictive regression model and shows that stock market *12-month return* (TMR), the time-series momentum predictor of Moskowitz, Ooi, and Pedersen (2012), forecasts the aggregate stock market negatively in good times and positively in bad times. The out-of-sample R^2s are 0.96% and 1.72% in good and bad times, or 1.28% and 1.41% in NBER economic expansions and recessions, respectively. The TMR predictability pattern holds in the cross-section of U.S. stocks and the international markets. Our study shows that the absence of return predictability in good times, an important finding of recent studies, is largely driven by the use of the popular one-state predictive regression model.

JEL Classification: C53; C58; G11; G12; G17

Keywords: Return predictability; good and bad times; 200-day moving average; business cycle

1 Introduction

Return predictability is widely considered a stylized fact. For example, Campbell and Thompson (2008), Cooper and Priestly (2009), Rapach, Strauss, and Zhou (2010), Binsbergen and Koijen (2010), Henkel, Martin, and Nardari (2011), Ferreira and Santa-Clara (2011), Dangl and Halling (2012), Li and Yu (2012), Li, Ng, and Swaminathan (2013), Pettenuzzo, Timmermann, and Valkanov (2014), Rapach, Ringgenberg, and Zhou (2016), and Møller and Rangvid (2014, 2017), among many others, provide empirical evidence that stock returns can be significantly predicted in- and out-of-sample. However, Rapach, Strauss, and Zhou (2010), Henkel, Martin, and Nardari (2011), and Dangl and Halling (2012) find a puzzling result that return predictability exists only in bad times (e.g., NBER-dated economic recessions). Theoretically, Cujean and Hasler (2017) provide a general equilibrium model to rationalize the asymmetric pattern of return predictability in good vs bad times. An open empirical question is still whether stock returns are really unpredictable at all in good times (out-of-sample).

This paper addresses the no predictability in good times puzzle and shows that stock returns can in fact be significantly predicted out-of-sample in good and bad times. The key idea is to use a two-state predictive regression model, instead of the traditional one-state predictive regression. As it turns out, the one-state predictive regression is misspecified for many existing predictors. With the two-state model, for example, the slope of the conditional market variance, one of the best economic predictors, changes signs from positive in good times to negative in bad times, and the net equity expansion, default yield spread, and inflation indicate predictability in good times while none in bad times. Therefore, the popular one-state predictive regression model is likely to understate the true degree of predictability in good times.

The two-state predictive regression model is similar to Cooper, Gutierrez, and Hameed (2004) and Boyd, Hu, and Jagannathan (2005) in other contexts. Since the NBER-dated business cycles are determined ex post and cannot be used for real time forecasting, we use the 200-day moving average (MA) indicator, which equals one (a good time) if the S&P 500 index level is above its 200-day MA and zero (a bad time) otherwise. According to Siegel (1994), the use of moving averages goes back at least to the 1930s. In practice, the 200-day MA has been widely used for identifying up and down markets, which we call good and bad times, and has been plotted for years in investment letters, trading softwares, and newspapers (e.g., Investor Business Daily). Intuitively, good times are those periods during which the stock market price is generally going up so that its current price is above its mean in the past 200 days, in response to perhaps a positive economic outlook. However, when the market price breaks down the 200-day MA, the economy may have fundamentally worsened or is expected to worsen.

There are two reasons to use the above simple ex ante two-state model. First, as argued by Cooper, Gutierrez, and Hameed (2004) and Boyd, Hu, and Jagannathan (2005), economic interpretations of the states are much easier to understand than a complex latent state econometric model. Indeed, our indicator is intuitive and is likely what investors may actually use to assess the market state in practice (Smith, Wang, Wang, and Zychowicz, 2016). Second, complex models are known to be counter-productive in out-of-sample forecasting. The more complex the model, the larger the estimation errors that worsen its performance. For instance, Lettau and Van Nieuwerburgh (2008) find that regime-shifting models perform poorly out-of-sample because of unreliable estimates of both the regime timing and size of the regime shifts.

How do the market states relate to macroeconomic business cycles? Over the sample period 1957:01–2015:12 (708 months), the 200-day MA indicator identifies 484 months as good times and the remaining 224 months as bad times. The proportion of bad times is 32%, close to the findings of Perez-Quiros and Timmermann (2000) and Henkel, Martin, and Nardari (2001) who identify about 30% of periods as bad times using sophisticated Bayesian learning approaches. Also, these bad times cover about 81% of NBER recessions and 86% of financial tressed periods (based on the Federal Reserve Bank of Kansas City Financial Stress Index). The less than 100% coverage of NBER recessions may not be necessarily a problem for our market state indicator. For instance, May and June of 2009 are identified as NBER recessions, but are identified as good times with our market state indicator. In fact, the stock market started a new bull market from March 2009 (the bottom was made in February). In general, although our bad times are defined ex ante and computed in real time, they relate to NBER recessions well.

When applying the two-state predictive regression to the 14 economic variables in Welch and Goyal (2008),[1] six variables display significance in good times and two display significance in

[1] Since there are some outliers, we replace the sample variance in Welch and Goyal (2008) with the conditional variance proposed by Johnson (2017).

bad times for in-sample forecasting. Specifically, the conditional variance is the only one that significantly predicts the market in both good and bad times. Its forecasting sign is positive in good times and negative in bad times, consistent with Ghysels, Plazzi, and Valkanov (2016) that the risk-return trade-off holds over samples excluding extremely bad times like financial crises. Moreover, net equity expansion, default yield spread, and inflation are not significant predictors with the one-state regression, but they significantly predict the market in good times with the two-state regression. Out-of-sample, nine variables display better performance. For example, the out-of-sample R^2 of conditional variance improves from -0.54% with the one-state regression to 0.87% with the two-state regression. However, due to their substantial forecasting instability as emphasized by Welch and Goyal (2008), these economic variables do not show significant predictive power in good times even in the two-state model.

We then turn our attention to a new predictor, the past 12-month return (TMR) of the S&P 500 index in excess of the one-month T-bill rate. In a seminar study, Moskowitz, Ooi, and Pedersen (2012) show that TMR of each futures contract (e.g., equity index, currency, mommodity, or sovereign bond) positively forecasts its future one to 12 month returns in-sample. We extend Moskowitz, Ooi, and Pedersen (2012) to spot market with the S&P 500 index as the underlying asset and explore whether TMR has out-of-sample forecasting power for the stock market.

Theoretically, TMR has solid economic reasons for being a predictor. First, Cujean and Hasler (2017) show that, in a general equilibrium model with two types of investors, the stock market is jointly driven by economic fundamental and investors' disagreement. Economic fundamental is mean-reverting and dominant in good times, causing price mean-reversal, whereas investor disagreement is counter-cyclical and dominant in bad times, leading one type of investors to underreact to (good) news and therefore causing short-term price momentum. Second, based on the general equilibrium model with technical traders of Han, Zhou, and Zhu (2016), TMR predicts the market negatively in good times when there are too many trend followers chasing price trends in the market. In bad times, there are much fewer trend followers, and so TMR predicts the market positively. While the popular one-state predictive regression model clearly cannot capture the changing slope property, our proposed two-state model serves well for the purpose.

Empirically, we find evidence consistent with the theoretical implications of Cujean and Hasler (2017) and Han, Zhou, and Zhu (2016). Over the full sample period 1957:01–2015:12, the slope of TMR in the two-state regression is negative in good times and positive in bad times, -0.60 ($t = -2.81$) vs 0.71 ($t = 2.79$). With 15 years of in-sample parameter training, the out-of-sample R^2 is 0.96% in good times and 1.72% in bad times, which are statistically significant at the 5% level. This performance is robust over NBER-dated business cycles, with out-of-sample R^2 1.28% in expansions and 1.41% in recessions. Overall, we provide, for the first time, strong evidence that the stock market is predictable in good times, as well as in bad times. This finding is important for both academics and practitioners because there are much more good times than bad times in the stock market.[2]

The predictability of TMR is not only statistically significant, but also economically substantial. Consider a mean-variance investor with a risk aversion of three who allocates his wealth between the market portfolio and the one-month T-bill. Instead of using the historical mean return, he uses the expected return predicted by TMR in his asset allocation decisions. Then the investor can earn an extra certainty equivalent return (CER) of 1.19% (1.40%) per year in good times (NBER expansions), and of 6.91% (12.59%) per year in bad times (NBER recessions). Although the annualized CERs are much larger in bad times (recessions) than good times (expansions), the values in good times (expansions) are economically sizeable.[3] Moveover, since the market is more often in good times (expansions), the compounding effect of the extra returns is substantial.

The forecasting pattern of TMR also holds in the cross-section of U.S. stocks. Specifically, when the U.S. stocks are sorted by size, book-to-market ratio, and industry, TMR significantly predicts their returns, with negative slopes in good times and positive slopes in bad times. The in- and out-of-sample R^2s are significant for most of portfolios in good and bad times, as well as over the NBER business cycles.

The predictability of TMR is quite robust. Internationally, by using the U.S. market state, the U.S. TMR significantly forecasts seven out of 11 countries in good times and ten countries in bad times. This result is consistent with the finding of Rapach, Strauss, and Zhou (2013) that the U.S. market is leading non-U.S. industrialized countries.

[2]According to the NBER-dated business cycles over 1957:01–2015:12, about 86% of times (607 out of 708 months) are identified as economic expansions. Henkel, Martin, and Nardari (2011) identify 72% of times as good times in the U.S. financial market.

[3]The CER in good times can be easily doubled if we replace the one-month T-bill rate with the market long-term mean in defining TMR.

Finally, we examine whether the predictive ability of TMR steps from cash flows, or discount rates, or both. Empirically, TMR strongly forecasts dividend-price ratio (a proxy of discount rate) and stochastic discount factors, pointing to the discount rate channel. It also significantly forecasts dividend growth, industrial production growth, and the Chicago Fed National Activity Index (CFNAI) in bad times. This evidence suggests that the predictive ability of TMR comes from both cash flows and discount rates.

The rest of the paper is organized as follows. Section 2 shows that a two-state predictive regression can improve the in- and out-of-sample forecasting power of traditional economic variables. Section 3 shows that TMR significantly predicts the aggregate market returns, negatively in good times and positively in bad times. The forecasting pattern also applies to cross-sectional and international portfolios. Section 4 explores the source of predictability, which is followed by a brief conclusion in Section 5.

2 Return Predictability and Market State

Our analysis focuses on one-month forecasting horizon. The reason is threefold. First, return predictability with a short horizon is usually magnified at longer horizons (Cochrane, 2011). Second, Boundoukh, Richardson, and Whitelaw (2008) show that long-horizon predictability may result from highly correlated sampling errors. Third, and most importantly, the goal of this paper is to show that stock returns can be significantly predicted in good and bad times, or in economic expansions and recessions. The NBER-dated economic recessions covered in this paper range from five to 17 months and longer horizon regressions would include random combinations of economic expansionary and recessionary periods that would undoubtedly blur the importance of the market states.

The target return to predict is the market return, which is defined as the log return on the S&P 500 index (including dividends) minus the log return on the one-month T-bill rate, available from the website of Amit Goyal. Due to some analysis using 30 years of market returns, the final sample period is 1957:01–2015:12, covering 708 months.

2.1 Market return in good and bad times

The market state is defined by the 200-day MA indicator, which equals one (a good time) if the S&P 500 index is above its 200-day MA, and zero otherwise (a bad time). As a real time indicator, the market state of next month is determined by the last trading day's 200-day MA indicator of current month. Statistically, one may search an optimal lag, say 100 days, to re-define the MA to yield better performance than what are reported in this paper. However, to mitigate the concerns of data mining and data snooping (see., e.g., Lo and MacKinlay, 1990), we simply use the well known 200-day MA that had been used by practitioners for decades before our out-of-sample period (see, e.g., Siegel, 1994), and that has been widely plotted in investment letters, trading softwares, and newspapers. Economically, since the 200-day MA is widely followed, its effect might be easy to understand. If enough investors believe it, they may herd on this information, thereby generating impact on stock price (see, e.g., Froot, Schaferstein, and Stein, 1992; and Bikhchandani, Hirshleifer, and Welch, 1992), making it interesting to examine whether it captures the good and bad times of the market.

Table 1 reports the summary statistics of the market return. Panel A assumes that the market always stays in one state. The average monthly return is 0.41% per month with a standard deviation of 4.24%. Consistent with the literature such as Albuquerque (2012), the skewness is negative, -0.69, suggesting that the market has more crashes than what a normally distributed market would experience. The market also displays weak time-series momentum with a positive first-order autocorrelation of 0.06. Moreover, the market stays much more often in macroeconomic expansions or in financial nonstressed states. Over 1957:01–2012:12, only 101 out of 708 months are identified as economic recessions according to the NBER-dated business cycles. Over 1990:02–2015:12, only 28 months are under financial market stress if a month is identified as stress if the Federal Reserve Bank of Kansas City Financial Stress Index (FSI) is one-standard deviation higher than its long-term mean.[4]

Panel B of Table 1 presents the results of assuming that the market switches between good and bad states, where the market state is defined by the 200-day MA indicator. In good times, the

[4]The FSI index is available at http://www.kc.frb.org/research/indicatorsdata/kcfsi/. It is the first principal component extracted from eleven financial variables that capture heightened stress in financial markets, including flight to quality/liquidity, increased uncertainty, and magnified information asymmetries.

average return, the kurtosis, and the realized Sharpe ratio are high; the volatility is low and the skewness is dramatically negative. In bad times, however, the average return, the skewness, and the realized Sharpe ratio are close to zero with a slightly negative sign; the volatility is high and the kurtosis is low. Overall, the bad times seem indeed bad. In addition, the first-order autocorrelation is negative in good times and positive in bad times, suggesting that the stock market reveals short-term time-series mean reversion in good times and momentum in bad times.

Panel B also shows that the 200-day MA indicator links well to business cycles and financial market stress conditions. It identifies 484 months as good times and the remaining 224 months as bad times, among which, bad times cover 81% of NBER recessions (82 out of 101 months), or 86% of FSI market stress periods (24 out of 28 months). This result is remarkable given that bad times are defined ex ante and computed in real time. It also should be mentioned that this less than 100% coverage of the NBER recessions is not necessarily a problem to the 200-day MA indicator. For example, May and June of 2009 are identified as NBER economic recessions, but are identified as good times by the 200-day MA indicator. In fact, the stock market started a new bull market from March 2009 (the bottom was February). Another example is October 1987, which is identified as an expansion by NBER but as a bad time by the 200-day MA indicator.

Finally, the 200-day MA indicator identifies 32% of months as bad times, which is much higher than the proportion of NBER recessions, 14%. This result is consistent with Perez-Quiros and Timmermann (2000) and Henkel, Martin, and Nardari (2001) who identify about 30% of periods as bad times with sophisticated Bayesian learning approaches. According to Paul A. Samuelson's quip that "Wall Street indexes predicted nine of the last five recessions!" there are more bad times or market crashes in the financial market than the real sector.[5]

2.2 Two-state regression

The standard one-state predictive regression model of forecasting the market return is

$$r_{t+1} = \alpha + \beta \cdot z_t + \varepsilon_{t+1}, \tag{1}$$

[5]From Paul A. Samuelson's "Science and Stocks" column in Newsweek, September 19th, 1966.

where z_t is the predictor. The forecasting power is based on the regression slope β or the R^2 statistic. If β is significantly different from zero or the R^2 is significantly larger than zero, one can conclude that z_t is a predictor of the market return. The out-of-sample forecast of next period's expected market return is computed recursively from

$$\hat{r}_{t+1} = \hat{\alpha}_t + \hat{\beta}_t \cdot z_t, \tag{2}$$

where $\hat{\alpha}_t$ and $\hat{\beta}_t$ are the ordinary least squares (OLS) estimates of α and β, respectively, based on data from the start of the available sample through t. The in-sample forecast is computed the same as above except that $\hat{\alpha}_t$ and $\hat{\beta}_t$ are replaced by those estimated by using the entire sample.

In this paper we follow Boyd, Hu, and Jagannathan (2005) and extend (1) by allowing for two market states: good and bad. That is, we run a two-state predictive regression as follows:

$$r_{t+1} = \alpha + \beta_{\text{good}} \cdot I_{\text{good},t} \cdot z_t + \beta_{\text{bad}} \cdot (1 - I_{\text{good},t}) \cdot z_t + \varepsilon_{t+1}, \tag{3}$$

where I_{good} is the 200-day MA indicator. Regression (3) nests regression (1) as a special case where the forecasts of the expected market return are the same in good and bad times.

We use Campbell and Thompson's (2008) out-of-sample R^2 statistic as the out-of-sample performance criterion, and define it as:

$$R^2_{OS} = 1 - \frac{\sum_{t=1}^{T}(r_t - \hat{r}_t)^2}{\sum_{t=1}^{T}(r_t - \bar{r}_t)^2}, \tag{4}$$

where T is the number of out-of-sample observations. $\hat{r}_t$ is the excess return forecast with (1) or (3), and $\bar{r}_t$ is the historical mean, both of which are estimated using data up to month $t-1$. If z_t is viable, R^2_{OS} will be positive and its mean-squared forecast error (hereafter MSFE) is lower than the MSFE with the forecast based on the historical mean. Campbell and Thompson (2008) show that a monthly R^2_{OS} of 0.5% can generate significant economic value. The null hypothesis of interest is therefore $R^2_{OS} \leq 0$ against the alternative hypothesis that $R^2_{OS} > 0$. We test this hypothesis by using the Clark and West (2007) MSFE-adjusted statistic. Define

$$f_{t+1} = (r_{t+1} - \bar{r}_{t+1})^2 - [(r_{t+1} - \hat{r}_{t+1})^2 - (\bar{r}_{t+1} - \hat{r}_{t+1})^2]. \tag{5}$$

The Clark and West (2007) MSFE-adjusted statistic is then the t-statistic from regressing f_{t+1} on a constant.

To investigate the asymmetry of return predictability, we separately compute the in- and out-of-sample R^2 statistics in good times (R^2_{good}) and in bad times (R^2_{bad}) as:

$$R^2_S = 1 - \frac{\sum_{t=1}^{T} I_{S,t}(r_t - \hat{r}_t)^2}{\sum_{t=1}^{T} I_{S,t}(r_t - \bar{r})^2}, \quad S = \text{good, bad}, \tag{6}$$

where $I_S = I_{\text{good}}$ (I_{bad}) equals one when the market state is good (bad) and zero otherwise. $\bar{r}$ is the return mean with the entire sample when calculating the in-sample statistics and is the return mean with data up to month $t-1$ when calculating the out-of-sample statistics. Note that, unlike the overall R^2 statistic, R^2_{good} and R^2_{bad} can be positive and negative. When necessary, we also calculate the in- and out-of-sample R^2_{exp} and R^2_{rec} statistics over NBER-dated economic expansions and recessions in the same way.

2.3 Forecasting the market with economic predictors

This section shows that the two-state predictive regression improves the forecasting performance of existing economic predictors. Table 2 considers the 14 well-known economic variables in Welch and Goyal (2008). Since there are some outliers, we replace the sample variance in Welch and Goyal (2008) with conditional variance, which is the fitted value from the regression of realized variance in month t on lagged realized variances in months $t-1$ and $t-12$ (Johnson, 2017). Table 2 presents the regression slope, t-statistic, in-sample R^2, and out-of-sample R^2_{OS}. Since most of these predictors are highly persistent, we use Newey-West t-statistics controlling for heteroskedasticity and autocorrelation with a lag of 12 throughout the paper.

Columns 2–4 of Table 2 reports the one-state regression results, which serves as a benchmark for comparison later. Among the 14 variables, four show significant in-sample forecasting power at the 10% level or stronger, including conditional variance, T-bill rate, long-term return, and term spread. The R^2 ranges from 0 for book-to-market ratio to 0.70% for long-term return, none of which exceeds 1%. We then use the first 15 years of data (180 months) for in-sample training and the remaining data for out-of-sample evaluation. As such, the out-of-sample period starts in

January 1972 and ends in December 2015. We use the expanding window approach to estimate the expected market return and restrict it to be non-negative as Campbell and Thompson (2008). The results show that three variables (i.e., long-term yield, long-term return, and inflation) generate non-negative R^2_{OS}s but none of them is statistically significant. This evidence confirms Welch and Goyal (2008) and many others that most, if not all, economic variables cannot predict the market return out-of-sample.

The rest of Table 2 then reports the results with two-state regression (3), where the market state is determined by the 200-day MA indicator. There are several interesting observations. First, the market can be significantly predicted in good times in-sample. Among the 14 variables, six reveal significance at the 10% level or stronger, including conditional variance, net equity expansion, T-bill rate, term spread, default yield spread, and inflation. Second, the predictability is asymmetric. Among the four variables that significantly predict the market with one-state regression, conditional variance is the best predictor that predicts the market positively in good times and negatively in bad times; T-bill rate and term spread predict the market only in good times, and long-term return predicts the market only in bad times. Third, net equity expansion, default yield spread, and inflation are not significant predictors with the one-state regression, but are significant predictors in good times with the two-state regression. This result suggests that the traditional one-state regression may underestimate the true degree of predictability in good times. Finally, the two-state regression significantly improves the forecasting power for many economic variables in terms of R-squares. For example, the R^2 and R^2_{OS} of conditional variance improve from 0.67% and -0.54% with one-state regression to 1.59% and 0.87% with two-state regression. To have an intuitive understanding, the last two columns of Table 2 report the R^2 and R^2_{OS} differences between the one- and two-state regressions. The results show that 12 variables display higher in-sample R^2s and nine variables display higher out-of-sample R^2_{OS}s.

In short, the two-state regression generally improve the forecasting power of existing economic predictors. However, as emphasized by Welch and Goyal (2008), the predictability is very unstable that leads a lack of predictive power out-of-sample. Unreported results show that there is still an absence of predictability for these predictors despite of the use of the two-state regression model. Hence, alternative predictors are needed.

3 Forecasting the market with TMR

In this section, we show that Moskowitz, Ooi, and Pedersen's (2012) time-series momentum predictor, TMR, fully resolves the no predictability puzzle based on our proposed two-state predictive regression (3).

3.1 TMR and its variants

In a seminal study, Moskowitz, Ooi, and Pedersen (2012) show that the past 12-month excess return, TMR, of each futures contract (out of 58) is a positive predictor of its future returns. This predictive effect persists for about one year and then partially reverses over longer horizons. They argue that this evidence is consistent but beyond asset pricing theories, such as initial under-reaction or delayed over-reaction.

Different from Moskowitz, Ooi, and Pedersen (2012) who use pooled regressions to incorporate information from other asset classes, we consider the standard time-series predictive regression for a given set of predictors, and assume that the market may switch between two states, good or bad. For robustness, we also consider three alternative definitions of TMR: 1) annualized volatility-scaled TMR, 2) TMR convergence divergence (TMRCD), i.e., the difference between TMR and its 30-year mean, and 3) annualized volatility-scaled TMRCD, where the volatility is the annualized moving standard deviation estimator as in Mele (2007). Volatility-scaled TMR captures the volatility timing ability (Fleming, Kirby, Ostdiek, 2001; Barroso and Santa-Clara, 2015; Moreira and Muir, 2017). TMRCD is to capture the property of reversion to long-term mean (see, e.g., Fama and French, 1988). This measure is also endorsed by practitioners. For example, John Bogle (2012), the legendary investor and founder of the Vanguard Group that manages billions of retirement funds for teachers and college professors, says that the #1 rule of investing (out of his ten rules) is "remember reversion to the mean."

Table 3 reports the results of predicting the market return with TMR, which is the main finding in this paper. Panel A shows that TMR does not do a good job with one-state regression (1). It positively forecasts the next-month market return but the predictability is not significant with a t-statistic of 0.64. The in-sample R^2 is as small as 0.09% and the out-of-sample R^2_{OS} is -0.11%, suggesting that TMR under-performs the historical mean in predicting the market return out-of-sample. The alternative TMRs fail to predict the market return either, with similar t-statistics and R^2s.

Panels B reports the results with two-state predictive regression (3), where the market state is determined by the 200-day MA indicator. In stark contrast to Panel A, TMR significantly forecasts the market return negatively in good times and positively in bad times. This forecasting pattern suggests that the stock market shows short-term time-series mean reversion in good times and momentum in bad times, and therefore, one-state regression under-estimates the TMR forecasting power. In good times, one standard deviation increase in TMR predicts a decrease of 0.60% in the expected market return. In contrast, in bad times, one standard deviation increase in TMR predicts an increase of 0.71% in the expected market return. The in-sample R^2 is 1.59% and the out-of-sample R^2_{OS} is 1.32%. As Campbell and Thompson (2008) show that even an R^2_{OS} of 0.5% yields significant economic value, the TMR predictive ability is apparently substantial and convincing.

Welch and Goyal (2008) provide an insightful graphical approach to evaluate the out-of-sample predictive power. This device depicts recursively the residuals to show whether the predictive regression forecast has a lower MSFE than the historical mean for any period by simply comparing the height of the curve at the beginning and end points of the segment corresponding to the period of interest. If the curve is higher at the end of the segment relative to the beginning, the predictive forecast has a lower MSFE during the period. A predictive forecast that always outperforms the historical mean will have a slope that is positive everywhere.

Figure 1 plots the cumulative sum-squared error (CSFE) from the historical mean forecast minus the CSFE from the competing forecast. The positive slopes reveal that TMR outperforms the historical mean forecast consistently over time. This is in contrast to Welch and Goyal (2008) who find that the out-of-sample predictive ability of the economic variables in Table 2 deteriorates remarkably after the oil shock of the mid-1970s. Our out-of-sample period, 1972:01–2015:12, allows us to analyze how TMR behaves over the recent market period characterized by the collapse of the technology bubble and the 2008–2009 mortgage crisis.

Table 3 also reports the in-sample R^2 and out-of-sample R^2_{OS} in good and bad times, as well as over NBER expansions and recessions, where good times are defined by the 200-day MA indicator.

The analysis also provides a hedge against the risk of data mining. With one-state regression, the R^2s and R^2_{OS}s of TMR and its alternatives are negative in good times and NBER expansions, and positive in bad times and recessions. This evidence is consistent with the literature and suggests that TMR can only predict the market in some special time periods, but not overall. With two-state regression, however, TMR and its alternatives generate positive R^2s and R^2_{OS}s in both good and bad times, and over NBER expansions and recessions. Moreover, with one exception, all these statistics are larger than 1%. Specifically, the TMR R^2 and R^2_{OS} are 1.98% and 0.96% in good times, and are 1.18% and 1.72% in bad times, respectively. Similarly, the R^2 and R^2_{OS} are 1.39% and 1.28% in NBER expansions, and they are 2.07% and 1.41% in NBER recessions, respectively.

To further reduce the concern that the outperformance of TMR is driven by certain occasional periods, Figure 2 plots the regression slopes over the out-of-sample period, 1972:01–2015:12. Apparently, the slopes in NBER-dated expansions and recessions are always different from zero, suggesting the robustness of TMR in predicting the market return. In sum, TMR consistently predicts the market and its performance exists in different market states and business cycles. This finding is important for investors since the market is more often in good times than in bad times.

3.2 Theoretical explanations

From an economic point of view, there are two possible explanations why TMR can predict the stock market with different signs, negative in good times and positive in bad times.

The first explanation is from Cujean and Hasler (2017). In a general equilibrium model, they consider an economy with two assets, a risk free asset and a risky asset (stock). The stock is a claim to its dividend process, where the expected dividend growth, henceforth the fundamental, is mean-reverting but unobservable. There are two types of investors, A and B. Both types of investors can observe the realized dividends, but have to learn the fundamental that affects the dividend they consume and the price of the stock they trade. Type A investors have the correct model that the fundamental evolves continuously over the business cycle, whereas type B investors use a discrete-state model by assuming that the fundamental switches between a good and a bad state. As such, type B investors face one more type of uncertainty, the economy state, in addition to the uncertainty of the fundamental.

Because the two types of investors assess uncertainty differently, they revise their expectations at different speeds and in different directions depending on the state of the economy. In good times, when a piece of good news occurs, type B investors only update their belief on the fundamental as the news is consistent with their belief on the economy state. In this case, the disagreement between the two types of investors is minimum. When a piece of bad news occurs, type B investors may think that the economy state has changed from good to bad and update their beliefs on the economy state and the fundamental, suggesting that Type B investors overreact to bad news and the disagreement increases. Therefore, the mean-reverting property of the fundamental and the overreaction of type B investors both suggest that the stock price display a mean-reverting property in good times. In bad times, however, good news is never good enough to convince type B investors that the economy is improving and their underreaction polarizes the disagreement, generating price momentum.

The second explanation is from Han, Zhou, and Zhu (2016). With a similar economy, they assume the type B investors in Cujean and Hasler (2017) are technical traders who infer information from the historical prices via moving averages. As a result, the price moving average serves as a state variable in the equilibrium price. When the ratio of technical traders are small, the usual trend following moving average strategy is profitable, which implies that there are not enough investors trading on this signal, therefore generating price momentum. When the ratio of technical traders are large, the trend following moving average strategy is unprofitable, which implies that there are too many investors trade on this signal, therefore generating price reversal.

3.3 Economic value

We now examine the economic value of TMR in predicting the market. Following Campbell and Thompson (2008) and Ferreira and Santa-Clara (2011), among others, we explore the annualized certainty equivalent return (CER) gain. The higher the CER gain, the larger the risk-rewarded return by using TMR.

Suppose a mean-variance investor invests his wealth between the market portfolio and the one-month T-bill rate. At the start of each month, he allocates a proportion of w_t to the market portfolio to maximize his next month' expected utility

$$U(R_p) = \mathrm{E}(R_p) - \frac{\gamma}{2}\mathrm{Var}(R_p), \tag{7}$$

where R_p is the (simple) return on the investor's portfolio, $\mathrm{E}(R_p)$ and $\mathrm{Var}(R_p)$ are the mean and variance of the portfolio return, and γ is the investor's coefficient of risk aversion.

Let r_{t+1} and $R_{f,t+1}$ be the excess return and T-bill rate. The investor's portfolio return at the end of each month is

$$R_{p,t+1} = w_t r_{t+1} + R_{f,t+1}, \tag{8}$$

where $R_{f,t+1}$ is known at t. With a simple calculation, the optimal portfolio is

$$w_t = \frac{1}{\gamma}\frac{\hat{r}_{t+1}}{\hat{\sigma}^2_{t+1}}, \tag{9}$$

where $\hat{r}_{t+1}$ and $\hat{\sigma}^2_{t+1}$ are the investor's estimates on the mean and variance of the market portfolio based on information up to time t.

The CER of the portfolio is

$$\mathrm{CER} = \hat{\mu}_p - \frac{\gamma}{2}\hat{\sigma}^2_p, \tag{10}$$

where $\hat{\mu}_p$ and $\hat{\sigma}^2_p$ are the mean and variance of the investor's portfolio over the out-of-sample evaluation period. The CER can be interpreted as the compensation to the investor for holding the market portfolio. The difference between the CERs for the investor using the predictive regression based on TMR and the historical mean as the forecast of the market return is naturally an economic measure of predictability significance.

Table 4 presents the welfare benefits generated by optimally trading on TMR for the investor with a risk aversion of three. That is, we report the CER difference between the strategy using the TMR forecast and the strategy using historical mean forecast of the market return. We annualize the CER by multiplying 1,200 so that the CER difference denotes the percentage gain per year for the investor to use the two-state regression forecast instead of the historical mean forecast. Following

Campbell and Thompson (2008), we assume that the investor uses a five-year moving window of past monthly returns to estimate the variance of the excess market return, and constraints w_t to lie between 0 and 1.5 to exclude short sales and to allow for at most 50% leverage.

In Panel A, when there is no transaction cost, the annualized CER gain by using the TMR is 2.96%, suggesting that investing with TMR forecast can generate 2.96% more risk-adjusted return relative to the historical mean forecast. In Panel B, when there is a transaction cost of 50 basis points, the CER gain by using TMR is 1.77%, which is still economically sizeable. Moreover, the CER gain can be easily improved if we use the alternative TMRs. For example, if we use volatility-scaled TMR, the portfolio gain increases to 4.06% without transaction costs and 3.00% with a transaction of 50 basis points. TMRCD and volatility-scaled TMRCD have even better performance regardless the existence of transaction costs.

We also report in Table 4 portfolio gains in good and bad times, and in NBER expansions and recessions separately. In so doing, we show that TMR can generate economic values in different market states. While the gain in bad times or in NBER recessions is larger, it is still economically large in good times or expansions. For example, even with 50 basis points of transaction cost, one can obtain 1.92% or 1.53% more CERs per year in good times or NBER expansions by using volatility-scaled TMRCD.

Figure 3 plots the cumulative wealth the mean-variance investor can obtain from using the one- and two-state predictive regressions (1) and (3), respectively. As a benchmark, we also plot the wealth with the historical mean forecast. We assume the investor has a risk aversion of three, begins with \$1, and reinvests all proceeds over the period January 1972 to December 2015. The figure shows that the terminal wealth in December 2015 is \$33 if the investor uses the historical mean forecast, \$41 if he uses the one-state predictive regression forecast, and \$106 if he uses the two-state predictive regression forecast. Apparently, the strategy with the two-state regression dramatically outperforms the other two strategies.

3.4 Forecasting cross-sectional portfolios

In previous sections, we show that TMR predicts the market statistically and economically, and the other three alternative TMRs reveal even stronger forecasting ability. As our goal in this paper is to show that stock returns are predictable in good times, rather than how to achieve the maximum predictability, we focus on TMR and only report its results hereafter to save the space.

Panel A of Table 5 presents the in- and out-of-sample forecasting results of TMR on size quintile portfolios. The regression slopes are negative in good times and positive in bad times, among which four are significant in good times and two are significant in bad times. The in- and out-of-sample R^2s are increasing in general. This finding is different from Ferson and Harvey (1991) that small firms are more predictable in that they are more affected by improving economic fundamentals but more vulnerable during economic downturns (Perez-Quiros and Timmermann, 2000). One reason for the increasing predictability is the increasing correlation between the size portfolios and the market portfolio, which ranges from 0.75 for the smallest size portfolio to 1 for the largest size portfolio. The returns we use in this paper are value-weighted and a larger size portfolio has a higher correlation with the market portfolio than a smaller size portfolio. As a result, TMR exhibits stronger predictive ability for portfolios with a larger weight in the market portfolio, since it significantly predicts the market. Also, the outperformance exists over both NBER expansions and recessions. For example, for the largest size portfolio, the R^2s are 1.41%, 1.35%, and 1.57%, and the R^2_{OS}s are 1.30%, 1.23%, and 1.46% over the whole period, expansions, and recessions, respectively.

Panel B shows that the TMR predictive ability applies to book-to-market quintile portfolios. Again, these five value portfolios are negatively predicted by TMR in good times and positively predicted in bad times. While the R^2s do not display a monotonic pattern, the R^2_{OS}s decrease in terms of the book-to-market ratio. That is, the lowest book-to-market (growth) portfolio is more predictable than the value portfolio. Their R^2_{OS}s are 1.26% and 0.24%, respectively. Over NBER business cycles, only the growth portfolio has significant R^2_{OS}s in both economic expansions and recessions.

Studies on industry portfolio predictability are relatively limited. Two exceptions are Cohen and Frazzini (2008) and Menzly and Ozbas (2010) who provide supporting evidence that some industry portfolios are predictable while others are not. However, these two papers focus on in-sample forecasting and do not explore the out-of-sample performance. Panel C of Table 5 consider the five industry portfolios from Kenneth French's data library. The results show that except HiTec, all of the other four industries are significantly predicted by TMR, in- and out-of-sample. For example, the R^2 and R^2_{OS} for the Cnsmr industry are 1.18% and 1.07%, respectively. Moreover, its R^2_{OS}s are similar in both expansions and recessions, 1.11% and 0.98%. The forecasting performance for Manuf is more pronounced. The R^2s over the sample period, expansions, and recessions are 1.58%, 1.44%, and 1.92%, and the corresponding R^2_{OS}s are 1.26%, 1.27%, and 1.25%, respectively.

In sum, when controlling for market states, TMR is a powerful predictor of cross-sectional portfolios. It forecasts future stock returns negatively in good times and positively in bad times.

3.5 Forecasting international stock markets

Rapach, Strauss, and Zhou (2013) show that the lagged U.S. market return significantly predicts returns in other non-U.S. industrialized countries in- and out-of-sample. So one natural question is whether the forecasting pattern of the U.S. TMR, negative in good times and positive in bad times, continues to exist in the international stock markets as well. We use MSCI total stock index returns at Datastream (denominated in U.S. dollars) to explore whether the U.S. TMR could predict other financial markets. We conduct the analysis from the perspective of U.S.-based investors, which is different from Rapach, Strauss, and Zhou (2013) who use local currencies. We calculate the log excess returns by using the one-month T-bill rate taken from the data library of Kenneth French. Following Rapach, Strauss, and Zhou (2013), we choose 11 countries, including Australia, Canada, France, Germany, Italy, Japan, Netherlands, Sweden, Switzerland, United Kingdom, and United States.

Table 6 reports the results. The sample period is from January 1970 to December 2015. For consistency, we choose the first 15 years' data as in-sample training and the out-of-sample performance is based on the period January 1985 to December 2015. In general, the lagged U.S. TMR significantly predicts most international markets and the forecasting power is stronger in bad times in terms of the regression slopes.[6] In good times, seven countries can be predicted, among which five are significant at the 5% level and two at the 10% level. In contrast, in bad times, except for Australia, all the other countries can be significantly predicted at the 10% level or stronger. The forecasting sign is consistent with the previous sections, negative in good times and positive in bad

[6]The stronger forecasting power in bad times can be partially related to the larger co-movements among stocks across countries during bad times than during good times indicated by many studies on asymmetric co-movements.

times. The R^2s range from 0.54% for Australia to 2.34% for U.S. Regarding out-of-sample, seven countries have significant R^2_{OS}s. We also report the forecasting results over NBER expansions and recessions. Over the out-of-sample period 1985–2015, five countries have significant R^2_{OS} in expansions, and five have significant R^2_{OS} in recessions. One natural improvement is to use country-specific good time indicators because different countries may have their own good and bad time cycles. We leave this for future research.

3.6 Forecasting the distribution of the market return

This section shows that TMR significantly predicts the distribution of the excess market returns, i.e., the left and right tails. We explore this problem for two reasons. First, as shown in Section 2, the relation between the future market return and TMR is nonlinear. The different patterns in predicting the left and right tails provide direct support for using the two-state regression. Second, forecasting crashes or downside risk is extremely important for risk management and asset allocation. In the literature, Cenesizoglu and Timmermann (2008) show that certain economic variables can predict one or two tails but not the center of stock returns. Baron and Xiong (2017) find that credit expansion predicts the left tail and the center, but not the right tail.

We use a quantile regression, which is robust to nonlinearality and outliers. Panel A of Table 7 considers the one-state quantile regression as:

$$Q_\tau(r_{t+1}|\mathrm{TMR}_t) = \alpha^{(\tau)} + \beta^{(\tau)}\mathrm{TMR}_t + \varepsilon_{t+1}, \quad \tau \in (0,1). \tag{11}$$

The results show that TMR significantly predicts both tails with different signs, but not the center of the distribution, which is consistent with our previous finding that, with the one-state predictive regression, TMR cannot predict the market. Specifically, TMR positively predicts the left tail and negatively predicts the right tail, suggesting that an increase in TMR is accompanied by a shrink of the distribution, i.e., the chance of large negative and positive returns decreases, whereas a decrease in TMR is accompanied by an increased dispersion in future market returns, i.e., the chance of large negative and positive returns increases.

Panel B of Table 7 presents the results with the following two-state quantile regression:

$$Q_\tau(r_{t+1}|\text{TMR}_t) = \alpha^{(\tau)} + \beta_{\text{good}}^{(\tau)} I_{\text{good},t} \cdot \text{TMR}_t + \beta_{\text{bad}}^{(\tau)} (1 - I_{\text{good},t}) \cdot \text{TMR}_t + \varepsilon_{t+1}, \tag{12}$$

where I_{good} is the 200-day MA indicator. Two observations follow the table immediately. First, with this two-state quantile regression, the median of the market return can be significantly predicted by TMR, negatively in good times and positively in bad times, confirming again that the market state is important in return predictability. Second, the predictability on the left and right tails exists only in bad times. In good times, TMR can only predicts market downside movements, and does not have any power for upside movements. Overall, if the one-state regression is true, the TMR slopes should be the same in two states. However, this table shows that the left tail slopes follow an asystematic pattern with large positive values in bad times and large negative values in good times, suggesting the necessity of controlling for the market state.

3.7 Robustness

As we define TMR with the past 12-month cumulative market returns, one interesting question is that whether the result still holds with different lags, rather than 12. Figure 4 plots the coefficients and t-statistics of the following regression for one to 12 lags,

$$r_{t+1} = \alpha + \beta_{\text{good}} \cdot I_{\text{good},t} \cdot \text{TMR}_{t-h,t} + \beta_{\text{bad}} \cdot (1 - I_{\text{good},t}) \cdot \text{TMR}_{t-h,t} + \varepsilon_{t+1}, \tag{13}$$

where $\text{TMR}_{t-h,t}$ is the past h-month cumulative excess market return. Apparently, the forecasting pattern holds in general and the forecasting power is increasing with respect to the month lag. Both good time and bad time slopes are significant when the lag increases to 10, which suggests that the lagged individual month returns contain complementary information on the future market returns.

Figure 4 also plots the R^2 and R^2_{OS} statistics. As the month lag increases from one to 12, the forecasting performance increases in general. This evidence is consistent with Geweke (1981) and Jegadeesh (1991) that using multi-month lagged returns to predict future returns can increase statistical power, and the choice of 12-month is not chosen by chance. Finally, untabulated results show that our finding continues to exist when we use 100- and 300-day MAs in defining the market state.

4 Forecasting Channel

Since a stock is priced by discounting its expected future cash flows with discount rates, the source of predictability is obviously from these two channels: cash flows or discount rates or both. This section attempts to explore the source.

According to the following decomposition as in Campbell and Shiller (1988) and Cochrane (2011),

$$\mathrm{DP}_t \approx r_{t+1} - \mathrm{DG}_{t+1} + \rho \cdot \mathrm{DP}_{t+1}, \tag{14}$$

where DP_t is the dividend-price ratio at time t, DG_{t+1} is the dividend growth rate at $t+1$, and ρ is the log-linearization constant. Cochrane (2011) argues that DP_{t+1} is a proxy of discount rates and DG_{t+1} is a proxy of cash flows. If a variable predicts future returns, it must predict either the discount rate (DP_{t+1}) or the cash flow (DG_{t+1}), or the both.

In addition to DP, we also consider stochastic discount rate factor (SDF) that can be directly extracted from portfolio returns. According to Cochrane (2005), the excess return of any asset is negatively related to its covariance with SDF,

$$\mathrm{E}_t(r_{t+1}) = -R_{f,t}\mathrm{Cov}_t(\mathrm{SDF}_{t+1}, r_{t+1}),$$

where SDF_{t+1} is a measure of the discount rate and $R_{f,t}$ is the gross risk-free rate at $t+1$. Projecting r_{t+1} on SDF_{t+1}, we have

$$r_{t+1} = a + b \cdot \mathrm{SDF}_{t+1} + \varepsilon_{t+1}.$$

If a variable can predict r_{t+1}, it must predict SDF_{t+1} or ε_{t+1}, and SDF_{t+1} captures the discount rates and ε_{t+1} measures the cash flows. Suppose there exist a risk-free asset and N risky assets in the market. A default and observable SDF_{t+1} is

$$\mathrm{SDF}_{t+1} = \frac{1}{R_{f,t}} + (1_N - \mu/R_{f,t})'\Sigma(R_{t+1} - \mu), \tag{15}$$

where μ and Σ are the mean and the covariance matrix of the N risky asset return R_{t+1}. In this paper, we consider three SDFs, which are extracted from 10 size portfolios, 10 value portfolios, and 10 industry portfolios, respectively.

In addition to DG, we consider monthly industrial production growth and Chicago Fed national Activity Index (CFNAI), which are widely used as the proxy of macroeconomic activity (Allen, Bali, and Tang, 2012).

Table 8 reports the results of the following regression as:

$$y_{t+1} = \alpha + \beta_{\text{good}} \cdot I_{\text{good},t} \cdot \text{TMR}_t + \beta_{\text{bad}} \cdot (1 - I_{\text{good},t}) \cdot r_{t-12,t} + \zeta y_t + \varepsilon_{t+1}, \tag{16}$$

where y_{t+1} is one of the discount rate or cash flow proxy and I_{good} is the 200-day MA indicator. Panel A presents the results of forecasting discount rates with TMR. The dividend-price ratio is significantly predicted by TMR,[7] positively in good times and negatively in bad times, which implies that an increase in TMR predicts an increase in DP_{t+1} in good times and a decrease in DP_{t+1} in bad times. Keeping the dividend constant, in good times, an increase in DP_{t+1} suggests that the price declines at $t+1$, i.e., the investment condition deteriorates and the stock market reverts, consistent with our finding that TMR negatively predicts the market in good times. In contrast, in bad times, a decrease in DP_{t+1} says that the stock price will go up and the investment setting is improving.

Similarly, TMR predicts SDFs positively in good times and negatively in bad times. Since SDF is contemporaneously negatively correlated with stock returns in general, the forecasting pattern of TMR on SDFs is consistent with that on dividend-price ratio. In good times, an increase in TMR predicts an increase in SDFs next period. In bad times, an increase in TMR predicts a decrease in SDFs.

Panel B of Table 8 presents the results of forecasting cash flows with TMR. In general, TMR positively predict cash flows, but is only significant in bad times. For dividend growth, the regression slope on TMR is 0.13 ($t = 1.54$) in good times and 0.20 ($t = 2.31$) in bad times, which suggest that one standard deviation increase in TMR predicts a 0.13% increase in good times and a 0.20% increase in bad times. For industrial production and CFNAI, the slopes on TMR are close

[7] The high in-sample R^2 is due to high persistence of the dividend-price ratio.

to zero in good times, but large in bad times. In sum, while the TMR forecasting ability is mainly from the discount rate channel, we cannot exclude the cash flow channel in general.

5 Conclusion

In this paper, we show that the predictability of stock returns may be understated due to the use of the popular one-state predictive regression model, which is likely misspecified for most predictors. In our proposed and easily implemented two-state regression model, the predictability of all existing predictors is generally improved. Moreover, we find that TMR, the time-series momentum predictor of Moskowitz, Ooi, and Pedersen (2012), significantly predicts the aggregate stock market, negatively in good times and positively in bad times. The predictability is of statistical and economic significance in both good and bad times, as well as during both NBER economic expansions and recessions. Theoretically, our empirical results are consistent with the general equilibrium models of Cujean and Hasler (2017) and Han, Zhou, and Zhu (2016).

There are a number of subjects that are of interest for future research. First, it will be straightforward to apply the two-state regression to other markets, such as commodity and currency markets, to see whether there are substantial improvements in the predictability as in the stock market. Second, to account for the dynamics of predictors, it will be valuable to extend our two-state regression into predictive systems along the line of Pástor and Stambaugh (2009). Finally, because stock returns are state dependent, it will be interesting to examine how corporate decisions are affected by market states.

References

Albuquerque, R.A., 2012. Skewness in stock returns: Reconciling the evidence on firm versus aggregate returns. Review of Financial Studies 25, 1630–1673.

Allen, L., Bali, T.G., Tang, Y., 2012. Does systemic risk in the financial sector predict future economic downturns? Review of Financial Studies 25, 3000–3036.

Baron, Matthew, and Wei Xiong, 2017, Credit expansion and neglected crash risk. Quarterly Journal of Economics 132, 713–764.

Barroso, P., Santa-Clara, P., 2015. Momentum has its moments. Journal of Financial Economics 116, 111–120.

Bikhchandani, S., Hirshleifer, D., Welch, I., 1992. A theory of fads, fashion, custom, and cultural change as informatioal cascades. Journal of Political Economy 100, 992–1026.

Binsbergen, J.H.v., Koijen, R., 2010. Predictive regressions: a present-value approach. Journal of Finance 65, 1439–1471.

Bogle, J.C., 2012. The Clash of the Cultures: Investment vs. Speculation. John Wiley & Sons.

Boudoukh, J., Richardson, M.P., Whitelaw, R.F., 2008. The myth of long-horizon predictability. Review of Financial Studies 21, 1577–1605.

Boyd, J., Hu, J., Jagannathan, R., 2005. The stock market's reaction to unemployment news: Why bad news is usually good for stocks. Journal of Finance 60, 649–672.

Campbell, J.Y., Shiller, R.J., 1988. Stock prices, earnings, and expected dividends. Journal of Finance 43, 661–676.

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.

Cenesizoglu, T., Timmermann, A., 2008. Is the distribution of stock returns predictable? Working paper, HEC Montreal and UCSD.

Clark, T.E., West, K.D., 2007. Approximately normal tests for equal predictive accuracy in nested models. Journal of Econometrics 138, 291–311.

Cochrane, J.H., 2005. Asset Pricing. Princeton, NJ.

Cochrane, J.H., 2011. Discount rates. Journal of Finance 66, 1047–1108.

Cohen, L., Frazzini, A., 2008. Economic links and predictable returns. Journal of Finance 63, 1977–2011.

Cooper, I., Priestly, R., 2009. Time-varying risk premiums and the output gap. Review of Financial Studies 22, 2801–2833.

Cooper, M.J., Gutierrez Jr., R.C., Hameed, A., 2004. Market states and momentum. Journal of Finance 59, 1345–1365.

Cujean, J., Hasler, M., 2017. Why does return predictability concentrate in bad times? Journal of Finance, forthcoming.

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

Fama, E., French, K., 1988. Permanent and termporary components of stock prices. Journal of Political Economy 96, 246–273.

Ferreira, M.I., Santa-Clara, P., 2011. Forecasting stock market returns: the sum of the parts is more than the whole. Journal of Financial Economics 100, 514–537.

Ferson, W., Harvey, C., 1991. The variation in economic risk premiums. Journal of Political Economy 99, 385–415.

Fleming, J., Kirby, C., Ostdiek, B., 2001. The economic value of volatility timing, Journal of Finance 56, 329–352.

Froot, K.A., Schaferstein, D.S., Stein, J.J., 1992. Herd on the street: Informational inefficiencies in a market with short-term speculation. Journal of Finance 47, 1461–1484.

Geweke, J., 1981. The approximate slopes of econometric tests. Econometrica 49, 1427–1442.

Ghysels, E., Plazzi, A., Valkanov, R., 2016. The risk-return relationship and financial crises. Working paper, UNC.

Han, Y., Zhou, G., Zhu, Y., 2016. A trend factor: Any economic gains from using information over investment horizons? Journal of Financial Economics 122, 352–375.

Henkel, S.J., Martin, J.S., Nadari, F., 2011. Time-varying short-horizon predictability. Journal of Financial Economics 99, 560–580.

Jegadeesh, N., 1991. Seasonality in stock price mean reversion: evidence from the U.S. and U.K. Journal of Finance 46, 1427–1444.

Johnson, T.L., 2017. Retur predictability revisited using weighted least squares. Working paper, University of Texas at Austin.

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

Li, J., Yu, J., 2012. Investor attention, psychological anchors, and stock return predictability. Journal of Financial Economics 104, 401–419.

Li, Y., Ng., D.T., Swaminathan, B., 2013. Predicting market returns using aggregate implied cost of capital. Journal of Financial Economics 104, 419–436.

Lo, A.W., MacKinlay, A.C., 1990. Data-snooping biases in tests of financial asset pricing models. Review of Financial Studies 3, 431–467.

Mele, A., 2007. Asymmetric stock market volatility and the cyclical behavior of expected returns. Journal of Financial Economics 86, 446–478.

Menzly, L., Ozbas, O., 2010. Market segmentation and the cross-predictability of returns. Journal of Finance 65, 1555–1580.

Møller, S.V., Rangvid, J., 2014. End-of-the-year economic growth and time-varying expected returns. Journal of Financial Economics 115, 136–154.

Møller, S.V., Rangvid, J., 2017. Global economic growth and expected returns around the world: The end-of-the-year effect. Management Science, forthcoming.

Moreira, A., Muir, T., 2017. Volatility-managed portfolios. Journal of Finance, forthcoming.

Moskowitz, T.J., Ooi, Y.H., Pedersen, L.H., 2012. Time series momentum. Journal of Financial Economics 104, 228–250.

Pástor, L., Stambaugh, R.F., 2009. Predictive systems: living with imperfect predictors. Journal of Finance 64, 1583–1628.

Perez-Quiros, G., Timmermann, A., 2000. Firm size and cyclical variations in stock returns. Journal of Finance 55, 1229–1262.

Pettenuzzo, D., Timmermann, A., Valkanov, R., 2014. Forecasting stock returns under economic constraints. Journal of Financial Economics 104, 517–553.

Rapach, D.E., Strauss, J.K., Zhou, G., 2010. Out-of-sample equity premium prediction: combination forecasts and links to the real economy. Review of Financial Studies 23, 821–862.

Rapach, D.E., Strauss, J.K., Zhou, G., 2013. International stock return predictability: What is the role of the United States? Journal of Finance 68, 1633–1662.

Rapach, D.E., Ringgenberg, M.C., Zhou, G., 2016. Short interest and aggregate stock returns. Journal of Financial Economics 121, 46–65.

Siegel, J.J., 1994. Stocks for the Long Run: A Guide to Selecting Markets for Long-Term Growth. Burr Ridge, Ill.: Irwin Professional Publishing.

Smith, D.M., Wang, N., Wang, Y., Zychowicz, E.J., 2016. Sentiment and the effectiveness of technical analysis: Evidence from the hedge fund industry. Journal of Financial and Quantitative Analysis 51, 1991–2013.

Welch, I., Goyal, A., 2008. A comprehensive look at the empirical performance of equity premium performance. Review of Financial Studies 21, 1455–1508.

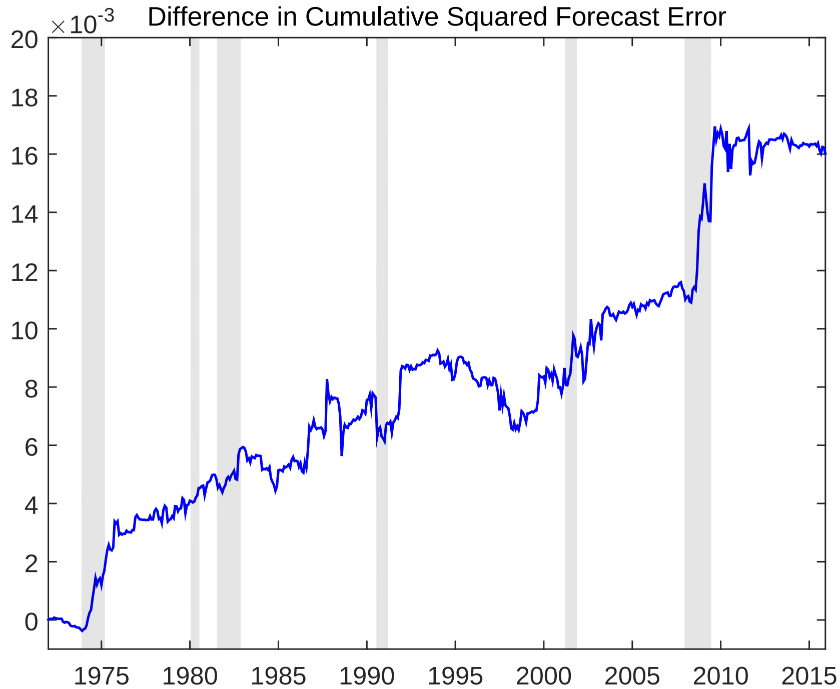

Figure 1 This figure depicts the difference between the cumulative squared forecast error (CSFE) of the historical mean forecast and the CSFE for the two-state predictive regression forecast with TMR, where TMR is the past 12-month market excess return. The whole sample period is 1957:01–2015:12 and the out-of-sample period is 1972:01–2015:12.

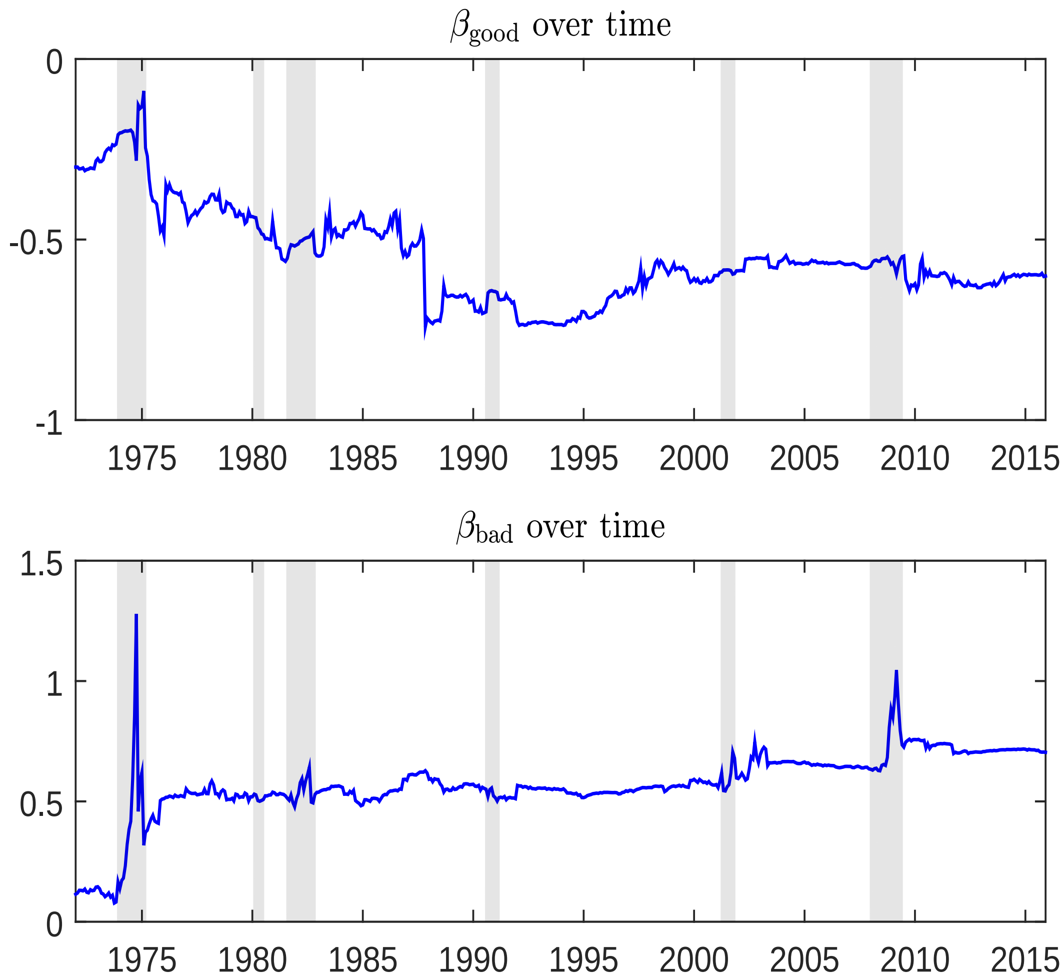

Figure 2 This figure plots the two-state regression slopes of forecasting the marekt excess return as:

$$r_{t+1} = \alpha + \beta_{\text{good}} \cdot I_{\text{good},t} \cdot \text{TMR}_t + \beta_{\text{bad}} \cdot (1 - I_{\text{good},t}) \cdot \text{TMR}_t + \varepsilon_{t+1},$$

where TMR is the past 12-month market exess return and I_{good} is the market state indicator that equals one when month t's last trading day close price of the S&P 500 index is above its 200-day moving average and zero otherwise. The whole sample period is 1957:01–2015:12 and the out-of-sample period is 1972:01–2015:12.

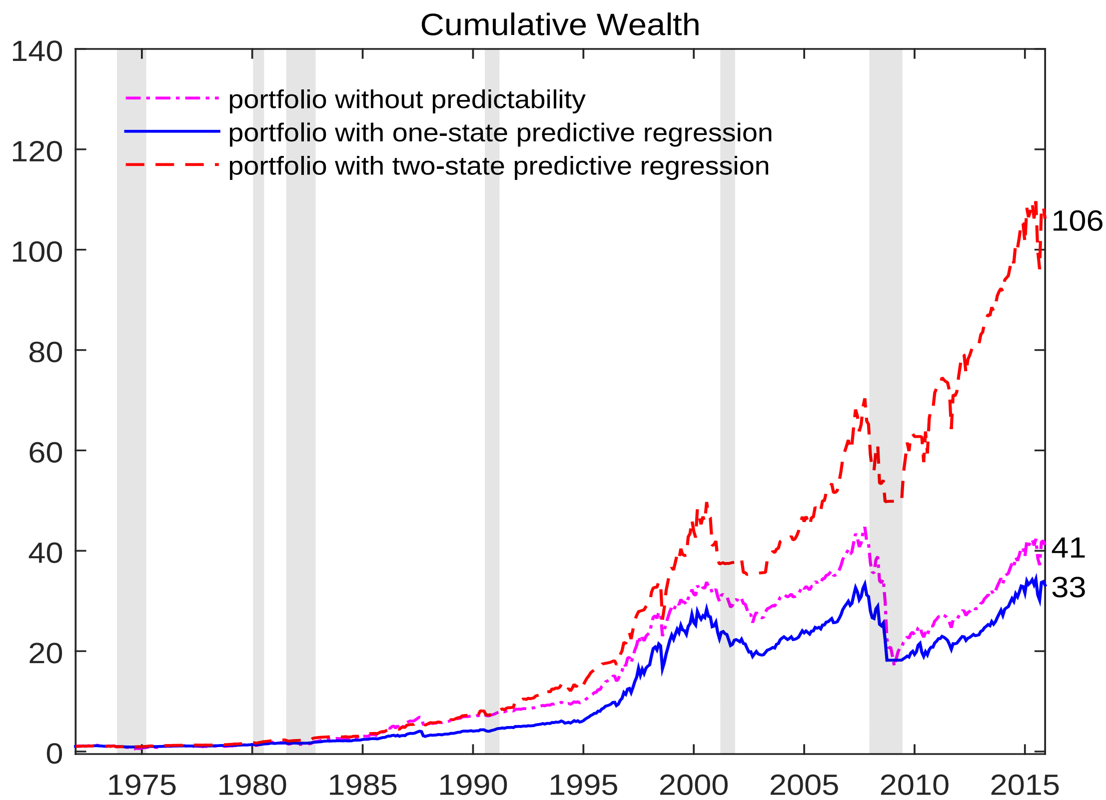

Figure 3 This figure plots the cumulative wealth for an investor who begins with $1 and reinvests all proceeds over time. The investor is equipped with a mean-variance preference and has a risk aversion of three. He allocates monthly between the S&P 500 index and one-month T-bill by assuming that the market excess return is unpredictable, or is predictable by the past 12-month market excess return, TMR, with one-state (one market state) regression or two-state (two market states) regression, where the market state is good if month t's last trading day close price of the S&P 500 index is above its 200-day moving average. The conditional variance of market return is estimated by the past five-year returns as Campbell and Thompson (2008). The whole sample period is 1957:01–2015:12 and the investment period is 1972:01–2015:12.

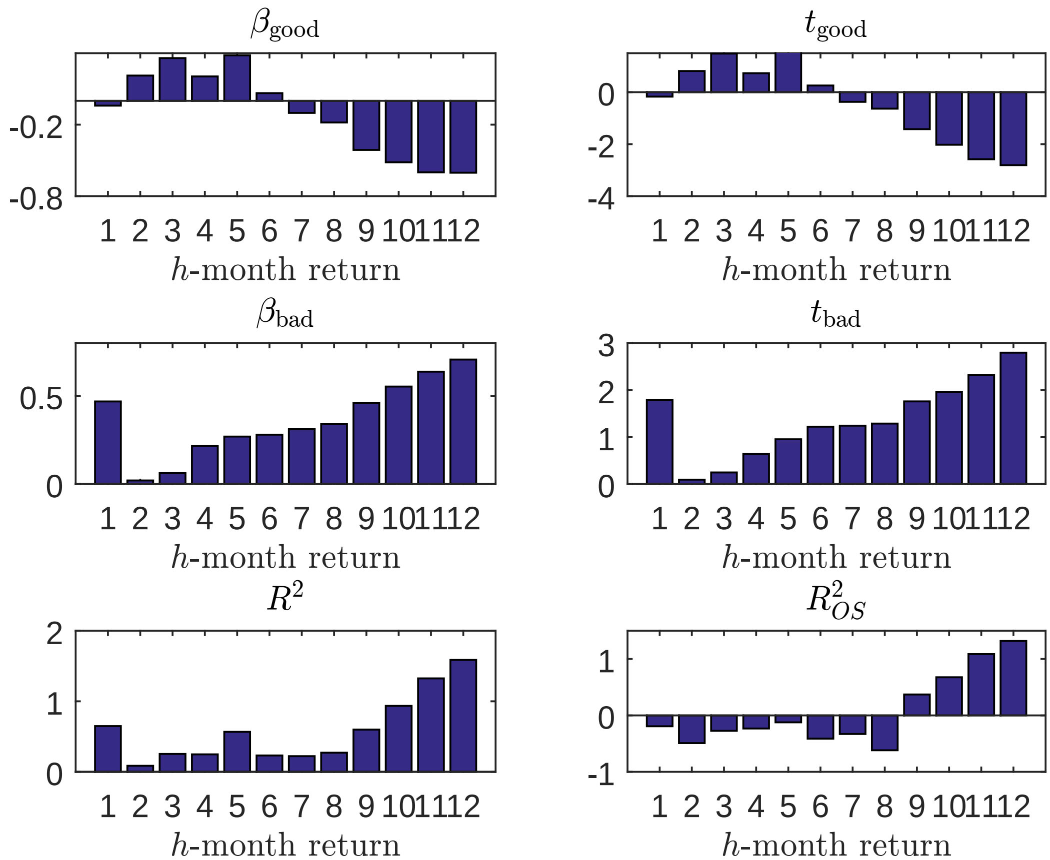

Figure 4 This figure plots the regression slopes, in-sample R^2, and out-of-sample R^2_{OS}s in good and bad times. We run the two-state regression of the monthly market return, r_{t+1}, on its past h-month excess return for $h = 1, 2, \cdots, 12$ as:

$$r_{t+1} = \alpha + \beta_{\text{good}} \cdot I_{\text{good},t} \cdot r_{t-h,t} + \beta_{\text{bad}} \cdot (1 - I_{\text{good},t}) \cdot r_{t-h,t} + \varepsilon_{t+1},$$

where $r_{t-h,t}$ is the excess return between month $t-h$ and month t. I_{good} is the market state indicator that equals one when month t's last trading day close price of the S&P 500 index is above its 200-day moving average and zero otherwise. The whole sample period is 1957:01–2015:12 and the out-of-sample period is 1972:01–2015:12.

Table 1 Summary statistics of the market excess return

This table reports the summary statistics of the monthly market return (the log return on the S&P 500 index in excess of the one-month T-bill rate) with one or two market states over the sample period 1957:01–2015:12. One state means that the market always stays in one state (no structural breaks), whereas two states mean that the market moves between two states, good times and bad times. Month $t+1$ is in good times if the market state indicator, I_{good}, equals one when month t's last trading day close price of the S&P 500 index is above its 200-day moving average and bad times otherwise. The statistics include the time-series average (Mean), standard deviation (Std), skewness (Skew), kurtosis (Kurt), the first-order autocorrelation ($\rho(1)$), and the monthly Sharpe ratio, which is defined as the average market return divided by its standard deviation. N is the number of time periods. #(NBER recession) represents the number of NBER-dated economic recessions. #(FSI stress) is the number of the Federal Reserve Bank of Kansas City Financial Stress Index (FSI) is one-standard deviation higher than its long-term mean over the period of 1990:02–2015:12.

	Mean	Std	Skew	Kurt	$\rho(1)$	Sharpe ratio	N	#(NBER recession)	#(FSI stress)
Panel A: One-state market									
	0.41	4.24	−0.68	5.49	0.06	0.10	708	101	28
Panel B: Two-state market									
Good times	0.65	3.66	−1.07	8.37	−0.06	0.18	484	19	4
Bad times	−0.11	5.26	−0.19	3.13	0.06	−0.03	224	82	24

Table 2 Forecasting the market excess return with popular economic variables

This table reports the results of forecasting the market excess return with popular economic variables by using the standard one-state predictive regression or a two-state regression. The market state indicator, I_{good}, equals one when month t's last trading day close price of the S&P 500 index is above its 200-day moving average and zero otherwise. R^2 is the in-sample R-square over the period 1957:01–2015:12 and R^2_{OS} is the Campbell and Thompson (2008) out-of-sample R-square with the first 15 years as the initial training period and with 1972:01–2015:12 as the evaluation period. All t-statistics are the Newey-West t-statistics controlling for heteroskedasiticity and autocorrelation. Statistical significance for R^2_{OS} is based on the p-value of the Clark and West (2007) MSPE-adjusted statistic for testing $H_0: R^2_{OS} \leq 0$ against $H_A: R^2_{OS} > 0$. ***, **, and * indicate significance at the 1%, 5%, and 10% levels, respectively.

	One-state regression				Two-state regression						Improvement	
Predictor	β	t-stat	R^2	R^2_{OS}	β_{good}	t-stat	β_{bad}	t-stat	R^2	R^2_{OS}	R^2	R^2_{OS}
Dividend-price ratio	0.17	0.96	0.16	−1.28	0.03	0.20	0.39	1.12	0.32	−1.21	0.16	0.07
Dividend yield	0.20	1.10	0.21	−1.42	0.03	0.20	0.46	1.33	0.45	−1.18	0.24	0.24
Earnings-price ratio	0.10	0.44	0.05	−0.12	−0.02	−0.10	0.21	0.66	0.13	0.04	0.08	0.16
Dividend-earnings ratio	0.08	0.44	0.04	−0.32	0.08	0.37	0.09	0.32	0.04	0.10	0.00	0.41
Conditional variance	-0.35^{*}	−1.65	0.67	−0.54	0.69^{*}	1.90	-0.53^{***}	−3.26	1.59	0.87^{**}	0.92	1.41
Book-to-market ratio	0.03	0.16	0.00	−0.79	−0.06	−0.28	0.14	0.39	0.05	−0.84	0.05	−0.05
Net equity expansion	−0.10	−0.39	0.06	−0.88	-0.37^{**}	−2.58	0.23	0.51	0.55	−0.99	0.50	−0.12
T-bill rate	-0.28^{*}	−1.72	0.44	−0.07	-0.37^{**}	−2.39	−0.15	−0.41	0.51	0.00	0.06	0.06
Long-term yield	−0.15	−0.89	0.12	0.14	−0.19	−1.05	−0.07	−0.22	0.14	0.29	0.02	0.15
Long-term return	0.36^{**}	2.45	0.70	0.06	0.26	1.21	0.48^{**}	2.11	0.77	−0.11	0.07	−0.18
Term spread	0.33^{**}	1.97	0.61	−0.70	0.40^{***}	2.61	0.21	0.57	0.66	−0.57	0.05	0.13
Default yield spread	0.16	0.69	0.15	−0.68	0.38^{**}	2.06	−0.10	−0.29	0.47	0.29^{**}	0.32	0.97
Default return spread	0.28	1.41	0.43	−0.42	0.29	1.10	0.27	0.93	0.43	−0.98	0.00	−0.56
Inflation	−0.09	−0.42	0.05	0.02	-0.39^{***}	−2.59	0.25	0.61	0.60	−0.28	0.56	−0.29

Table 3 Forecasting the market excess return with state-independent and dependent regressions

This table reports the results of forecasting the excess market return with one- and two-state regressions, respectively. The one-state regression refers to

$$r_{t+1} = \alpha + \beta z_t + \varepsilon_{t+1},$$

and the two-state regression refers to

$$r_{t+1} = \alpha + \beta_{\text{good}} I_{\text{good},t} \cdot z_t + \beta_{\text{bad}}(1 - I_{\text{good},t}) \cdot z_t + \varepsilon_{t+1},$$

where z_t is the predictor, and I_{good} is the market state indicator that equals one when month t's last trading day close price of the S&P 500 index is above its 200-day moving average and zero otherwise. TMR is the 12-month excess return on the S&P 500 index, volatility-scaled TMR is the TMR scaled by its annualized volatilty, TMRCD is the 12-month market return convergence divergence and is defined as the difference between TMR and its 30-year mean, and volatility-scaled TMRCD is the TMRCD scaled by its annualized volatilty. R^2 is the in-sample R-square over the period 1957:01–2015:12 and R^2_{OS} is the Campbell and Thompson (2008) out-of-sample R-square over 1972:01–2015:12. R^2_{good} (R^2_{bad}) is calculated seperately in good (bad) times, and R^2_{exp} (R^2_{rec}) is calculated over NBER-dated economic expansions (recessions), respectively. All t-statistics are the Newey-West t-statistic controlling for heteroskedasiticity and autocorrelation. Statistical significance for R^2_{OS} is based on the p-value of the Clark and West (2007) MSPE-adjusted statistic for testing $H_0 : R^2_{OS} \leq 0$ against $H_A : R^2_{OS} > 0$. ***, **, and * indicate significance at the 1%, 5%, and 10% levels, respectively.

Panel A: One-state regression

	Overall				Good vs bad times				Expansions vs recessions			
Predictor	β	t-stat	R^2	R^2_{OS}	R^2_{good}	R^2_{bad}	$R^2_{OS,\text{good}}$	$R^2_{OS,\text{bad}}$	R^2_{exp}	R^2_{rec}	$R^2_{OS,\text{exp}}$	$R^2_{OS,\text{rec}}$
TMR	0.13	0.64	0.09	−0.11	−0.49	0.69	−0.48	0.30	−0.18	0.76	−0.38	0.54
Volatility-scaled TMR	0.15	0.80	0.13	−0.05	−0.65	0.95	−0.56	0.50	−0.30	1.21	−0.38	0.76
TMRCD	0.13	0.66	0.10	−0.12	−0.52	0.75	−0.59	0.40	−0.19	0.81	−0.44	0.66
Volatility-scaled TMRCD	0.21	1.04	0.24	0.00	−0.82	1.34	−0.87	0.96^{**}	−0.34	1.66	−0.48	1.18^{*}

Panel B: Two-state regression

	Overall						Good vs bad times				Expansions vs recessions			
Predictor	β_{good}	t-stat	β_{bad}	t-stat	R^2	R^2_{OS}	R^2_{good}	R^2_{bad}	$R^2_{OS,\text{good}}$	$R^2_{OS,\text{bad}}$	R^2_{exp}	R^2_{rec}	$R^2_{OS,\text{exp}}$	$R^2_{OS,\text{rec}}$
TMR	−0.60	−2.81	0.71	2.79	1.59	1.32^{***}	1.98	1.18	0.96^{**}	1.72^{**}	1.39	2.07	1.28^{**}	1.41^{**}
Volatility-scaled TMR	−0.57	−2.89	0.97	3.21	2.15	1.43^{***}	1.86	2.46	1.14^{**}	1.74^{**}	1.15	4.63	1.36^{**}	1.60^{**}
TMRCD	−0.57	−2.66	0.72	2.70	1.60	1.80^{***}	1.85	1.34	1.72^{***}	1.89^{**}	1.48	1.90	1.70^{***}	2.04^{**}
Volatility-scaled TMRCD	−0.49	−2.45	1.03	3.27	2.26	1.71^{***}	1.44	3.12	1.44^{***}	2.02^{***}	1.16	5.02	1.58^{***}	2.05^{**}

Table 4 Portfolio gain from using two-state regression (in % per year)

This table reports the CER gain for a mean-variance investor with a risk-aversion of three, who allocates monthly between the aggregate market and the one-month T-bill with the two-state predictive regression as:

$$r_{t+1} = \alpha + \beta_{\text{good}} I_{\text{good},t} \cdot z_t + \beta_{\text{bad}}(1 - I_{\text{good},t}) \cdot z_t + \varepsilon_{t+1},$$

where z_t is the predictor, and I_{good} is the market state indicator that equals one when month t's last trading day close price of the S&P 500 index is above its 200-day moving average and zero otherwise. TMR is the 12-month excess return on the S&P 500 index, volatility-scaled TMR is the TMR scaled by its annualized volatilty, TMRCD is the 12-month market return convergence divergence and is defined as the difference between TMR and its 30-year mean, and volatility-scaled TMRCD is the TMRCD scaled by its annualized volatilty. CER gain (in % per year) is the certainty-equivalent return gain for the investor. The portfolio weights are estimated recursively and are constrained to lie between 0 and 1.5, using data available at the forecast formation time t. The conditional variance in forming month t's portfolio is estimated by the past five-year returns. The full sample period is 1957:01–2015:12 and the investment period is 1972:01–2015:12.

Predictor	Overall	Good times	Bad times	NBER expansions	NBER recessions
Panel A: with no transactoin cost					
TMR	2.96	1.19	6.91	1.40	12.59
Volatility-scaled TMR	4.06	1.79	9.16	2.29	15.08
TMRCD	4.37	2.57	8.42	2.41	16.55
Volatility-scaled TMRCD	4.63	2.65	9.07	2.64	16.97
Panel B: with 50 bps transactoin costs					
TMR	1.77	0.21	5.98	0.20	11.68
Volatility-scaled TMR	3.00	0.98	8.18	1.21	14.27
TMRCD	3.20	1.62	7.48	1.21	15.49
Volatility-scaled TMRCD	3.54	1.92	7.92	1.53	16.18

Table 5 Forecasting cross-sectional portfolios with two-state regression

This table reports the results of forecasting cross-sectional excess portfolio returns with two-state predictive regression as:

$$r^i_{t+1} = \alpha^i + \beta^i_{\text{good}} \cdot I_{\text{good},\, t} \cdot \text{TMR}_t + \beta^i_{\text{bad}} \cdot (1 - I_{\text{good},\, t}) \cdot \text{TMR}_t + \varepsilon^i_{t+1},$$

where r^i_{t+1} is the excess return of the ith size, book-to-market ratio, industry, or momentum portfolio in month $t+1$. TMR is the past 12-month makret excess return and I_{good} is the market state indicator that equals one when month t's last trading pday close price of the S&P 500 index is above its 200-day moving averge and zero otherwise. R^2 is the in-sample R-square over 1957:01–2015:12 and R^2_{OS} is the Campbell and Thompson (2008) out-of-sample R^2_{OS} over 1972:01–2015:12. R^2_{exp} (R^2_{rec}) statistics are calculated over NBER-dated expansions (recessions). Statistical significance for R^2_{OS} is based on the p-value for the Clark and West (2007) out-of-sample MSPE-adjusted statistic for testing $H_0 : R^2_{OS} \leq 0$ against $H_A : R^2_{OS} > 0$. ***, **, and * denote significance at the 1%, 5%, and 10% levels, respectively.

Portfolio	In-sample					Out-of-sample		
	β_{good}	β_{bad}	R^2	R^2_{exp}	R^2_{rec}	R^2_{OS}	$R^2_{OS,\text{exp}}$	$R^2_{OS,\text{rec}}$
Panel A: Size portfolios								
Small	−0.56	0.56	0.53	0.33	1.11	−0.07	−0.45	0.98
2	−0.69**	0.52	0.69	0.57	1.02	0.31	0.15	0.72
3	−0.64**	0.47	0.70	0.63	0.88	0.30	0.20	0.53
4	−0.74***	0.53*	1.04	1.11	0.87	0.73**	0.76**	0.65
Large	−0.59***	0.64***	1.41	1.35	1.57	1.30***	1.23**	1.46*
Panel B: Book-to-market portfolios								
Low	−0.65***	0.66**	1.21	1.09	1.52	1.26***	1.10**	1.66**
2	−0.74***	0.59**	1.42	1.40	1.46	1.11**	1.22**	0.79
3	−0.57***	0.48*	0.97	0.91	1.15	0.30	0.42	−0.02
4	−0.60***	0.65*	1.28	1.06	1.68	0.26	0.25	0.28
High	−0.57*	0.59*	0.84	0.83	0.86	0.24	0.01	0.81
Panel C: Five industry portfolios								
Cnsmr	−0.68***	0.55**	1.18	1.08	1.45	1.07**	1.11**	0.98
Manuf	−0.66***	0.67***	1.58	1.44	1.92	1.26**	1.27**	1.25*
HiTec	−0.37	0.49*	0.45	0.40	0.60	0.35	0.15	1.01
Hlth	−0.75***	0.56**	1.10	1.04	1.28	1.17**	1.30**	0.78
Other	−0.78***	0.66*	1.22	1.18	1.31	0.78**	0.77*	0.79

Table 6 Forecasting international excess market returns with two-state regression

This table reports the results of forecasting international excess market returns with two-state predictive regression as:

$$r^i_{t+1} \quad = \quad \alpha^i + \beta^i_{\text{good}} \cdot I_{\text{good},t} \cdot \text{TMR}_t + \beta^i_{\text{bad}} \cdot (1 - I_{\text{good},t}) \cdot \text{TMR}_t + \varepsilon_{t+1},$$

where r^i_{t+1} is the excess return of country i in month $t+1$, TMR is the past 12-month excess return of the S&P 500 index, and I_{good} is the market state indicator that equals one when month t's last trading day close price of the S&P 500 index is above its 200-day moving averge and zero otherwise. R^2 is the in-sample R-square over the sample period 1970:01–2015:12 and R^2_{OS} is the Campbell and Thompson (2008) out-of-sample R^2_{OS} over 1980:01–2015:12. All t-statistics are the Newey-West t-statistics controlling for heteroskedasiticity and autocorrelation. Statistical significance for R^2_{OS} is based on the p-value for the Clark and West (2007) out-of-sample MSPE-adjusted statistic for testing $H_0 : R^2_{OS} \leq 0$ against $H_A : R^2_{OS} > 0$. ***, **, and * denote significance at the 1%, 5%, and 10% levels, respectively.

	In-sample					Out-of-sample		
	β_{good}	β_{bad}	R^2	R^2_{exp}	R^2_{rec}	R^2_{OS}	$R^2_{OS,\text{exp}}$	$R^2_{OS,\text{rec}}$
Australia	−0.82	0.53	0.54	1.07	−0.78	−0.11	−0.38	0.89
Canada	−0.62**	0.68*	0.84	0.79	0.94	0.30	−0.41	1.72*
France	−1.09**	0.94**	1.50	1.84	0.41	0.71**	0.64*	0.92
Germany	−0.70*	0.96**	1.21	1.62	−0.02	0.56*	0.28	1.40
Italy	−0.81	1.38***	1.76	1.08	4.34	0.49*	0.40*	0.86
Japan	−0.26	1.16***	1.72	0.96	4.23	0.01	−1.21	4.09**
Netherlands	−0.67*	0.93**	1.45	2.13	−0.05	0.70*	0.28	1.74*
Sweden	−0.35	0.74*	0.56	0.90	−0.40	0.24	−0.25	1.68
Switzerland	−0.91***	1.04***	2.30	3.14	−0.17	1.65***	1.27**	2.94**
United Kingdom	−1.01***	1.11**	1.93	2.29	1.31	1.08***	0.72**	2.08**
United States	−0.82***	0.87***	2.34	2.42	2.17	1.40***	1.13**	2.12

Table 7 Forecasting the market return with quantile regression

Panel A considers state-independent τth quantile forecasting as

$$Q_\tau(r_{t+1}|\mathrm{TMR}_t) = \alpha^{(\tau)} + \beta^{(\tau)}\mathrm{TMR}_t.$$

Panels B reports the results on two-state τth quantile forecasting as

$$Q_\tau(r_{t+1}|\mathrm{TMR}_t) = \alpha^{(\tau)} + \beta^{(\tau)}_{\mathrm{good}} I_{\mathrm{good},t} \cdot \mathrm{TMR}_t + \beta^{(\tau)}_{\mathrm{bad}}(1 - I_{\mathrm{good},t}) \cdot \mathrm{TMR}_t,$$

where TMR is the past 12-month makret excess return and I_{good} is the market state indicator that equals one when month t's last trading day close price of the S&P 500 index is above its 200-day moving averge and zero otherwise. The sample period is 1957:01–2015:12.

Percentile	5	10	25	50	75	90	95
Panel A: One-state quantile regression							
β	1.35***	1.17***	0.45**	0.00	−0.54***	−0.68***	−0.69**
	(3.67)	(4.02)	(2.37)	(0.02)	(−3.83)	(−2.71)	(−2.11)
Panel B: One-state quantile regression							
β_{good}	−1.56**	−0.45	−0.67*	−0.46*	−0.42	0.10	−0.06
	(−2.48)	(−0.88)	(−1.76)	(−1.72)	(−1.50)	(0.30)	(−0.14)
β_{bad}	2.85***	2.68***	1.38***	0.65***	−0.65***	−1.05***	−1.42***
	(5.28)	(5.94)	(4.43)	(2.68)	(−3.06)	(−3.76)	(−2.96)

Table 8 Forecasting discount rates and cash flows

This table reports the results of forecasting discount rates and cash flows as

$$y_{t+1} = \alpha + \beta_{\text{good}} \cdot I_{\text{good},t} \cdot \text{TMR}_t + \beta_{\text{bad}} \cdot (1 - I_{\text{good},t}) \cdot \text{TMR}_t + \zeta y_t + \varepsilon_{t+1},$$

where TMR is the past 12-month makret excess return and I_{good} is the market state indicator that equals one when month t's last trading day close price of the S&P 500 index is above its 200-day moving averge and zero otherwise. SDF is the model-free and minimum variance stochastic discount factor constructed by the size, value, or industry portfolios, repectively. CFNAI is the Chicago Fed National Activity Index. All t-statistics are the Newey-West t-statistics controlling for heteroskedasiticity and autocorrelation. The sample period is 1957:01–2015:12.

y_{t+1}	β_{good}	t-stat	β_{bad}	t-stat	R^2
Panel A: Forecasting discount rates					
Dividend-price ratio	0.66^{***}	2.67	-0.55^{**}	-2.02	98.80
SDF with 10 size portfolios	2.82^{***}	4.06	-2.66^{**}	-2.02	1.89
SDF with 10 value portfolios	1.95^{*}	1.88	-2.35	-1.58	0.69
SDF with 10 industry portfolios	2.21^{*}	1.92	-3.65^{***}	-2.63	1.22
Panel B: Forecasting cash flows					
Dividend growth	0.13	1.54	0.20^{**}	2.31	8.57
Industrial production growth	0.03	0.83	0.23^{***}	5.88	4.53
CFNAI	0.01	0.26	0.21^{***}	4.15	4.34