Journal of Financial Economics 100 (2011) 514–537

Contents lists available at ScienceDirect

Journal of Financial Economics

journal homepage: www.elsevier.com/locate/jfec

Forecasting stock market returns: The sum of the parts is more than the whole ☆

Miguel A. Ferreira [a,b], Pedro Santa-Clara [a,c,*]

[a] Universidade Nova de Lisboa, Faculdade de Economia, Lisboa, Portugal
[b] European Corporate Governance Institute, Brussels, Belgium
[c] National Bureau of Economic Research, Cambridge, Massachusetts, USA

ARTICLE INFO

Article history:
Received 20 January 2010
Received in revised form
15 July 2010
Accepted 13 August 2010
Available online 19 February 2011

JEL classification:
G11
G12
G14

Keywords:
Predictability
Stock returns
Equity premium
Predictive regressions
Trading strategies

ABSTRACT

We propose forecasting separately the three components of stock market returns—the dividend–price ratio, earnings growth, and price–earnings ratio growth—the sum-of-the-parts (SOP) method. Our method exploits the different time series persistence of the components and obtains out-of-sample *R*-squares (compared with the historical mean) of more than 1.3% with monthly data and 13.4% with yearly data. This compares with typically negative *R*-squares obtained in a similar experiment with predictive regressions. The performance of the SOP method comes mainly from the dividend–price ratio and earnings growth components, and the robustness of the method is due to its low estimation error. An investor who timed the market using our method would have had a Sharpe ratio gain of 0.3.

1. Introduction

A long literature exists on forecasting stock market returns using price multiples, macroeconomic variables, corporate actions, and measures of risk.[1] These studies find evidence in favor of return predictability in sample. However, a number of authors question the findings on

☆ We thank an anonymous referee, Michael Brandt, Jules van Binsbergen, John Campbell, John Cochrane, Amit Goyal, Lubos Pastor, Bill Schwert (the editor), Ivo Welch, Motohiro Yogo, and Jialin Yu; seminar participants at Barclays Global Investors, Goethe Universität Frankfurt, Hong Kong University of Science and Technology Finance Symposium, Manchester Business School, BI-Norwegian School of Management, and University of Piraeus; and participants at the 2009 European Finance Association Conference, 2010 American Finance Association Conference, and CSEF-IGIER Symposium on Economics and Institutions for helpful comments. We are particularly grateful to Carolina Almeida, Filipe Lacerda, and Tymur Gabuniya for outstanding research assistance. This research is supported by a grant from the Fundação para a Ciência e Tecnologia (PTDC/EGE-GES/101414/2008).
* Corresponding author.
E-mail address: psc@fe.unl.pt (P. Santa-Clara).

0304-405X/$ - see front matter © 2011 Elsevier B.V. All rights reserved.
doi:10.1016/j.jfineco.2011.02.003

[1] Researchers who use the dividend yield include Dow (1920), Campbell (1987), Fama and French (1988), Hodrick (1992), Campbell and Yogo (2006), Ang and Bekaert (2007), Cochrane (2008), and Binsbergen and Koijen (2010). The earnings–price ratio is used by Campbell and Shiller (1988) and Lamont (1998). The book-to-market ratio is used by Kothari and Shanken (1997) and Pontiff and Schall (1998). The short-term interest rate is used by Fama and Schwert (1977), Campbell (1987), Breen, Glosten, and Jagannathan (1989), and Ang and Bekaert (2007). Inflation is used by Nelson (1976), Fama and Schwert (1977), Ritter and Warr (2002), and Campbell and Vuolteenaho (2004). The term and default yield spreads are used by Campbell (1987) and Fama and French (1988). The consumption–wealth ratio is used by Lettau and Ludvigson (2001). Corporate issuing activity is used by Baker and Wurgler (2000) and Boudoukh, Michaely, Richardson, and Roberts (2007). Stock volatility is used by French, Schwert, and Stambaugh (1987), Goyal and Santa-Clara (2003), Ghysels, Santa-Clara, and Valkanov (2005), and Guo (2006).

the grounds that the persistence of the forecasting variables and the correlation of their innovations with returns might bias the regression coefficients and affect t-statistics (Nelson and Kim, 1993; Cavanagh, Elliott, and Stock, 1995; Stambaugh, 1999; Lewellen, 2004). A further problem is the possibility of data mining illustrated by a long list of spurious predictive variables that regularly show up in the press, including hemlines, football results, and butter production in Bangladesh (Foster, Smith, and Whaley, 1997; Ferson, Sarkissian, and Simin, 2003). The predictability of stock market returns thus remains an open question.

In important recent research, Goyal and Welch (2008) examine the out-of-sample performance of a long list of predictors. They compare forecasts of returns at time $t+1$ from a predictive regression estimated using data up to time t with forecasts based on the historical mean in the same period. They find that the historical mean has better out-of-sample performance than the traditional predictive regressions. Goyal and Welch (2008) conclude that "these models would not have helped an investor with access only to available information to profitably time the market" (p. 1455) (see also Bossaerts and Hillion, 1999). While Inoue and Kilian (2004) and Cochrane (2008) argue that this is not evidence against predictability per se but only evidence of the difficulty in exploiting predictability with trading strategies, the Goyal and Welch (2008) challenge remains largely unanswered.

We offer an alternative approach to predict stock market returns: the sum-of-the-parts (SOP) method. We decompose the stock market return into three components—the dividend–price ratio, the earnings growth rate, and the price–earnings ratio growth rate—and forecast each component separately, exploiting their different time series characteristics. Because the dividend–price ratio is highly persistent, we forecast it using the currently observed dividend–price ratio. Because earnings growth is close to unpredictable in the short-run but has a low-frequency predictable component (Binsbergen and Koijen, 2010), we forecast it using its long-run historical average (20-year moving average). Finally, we assume no growth in the price–earnings ratio in this simplest version of the SOP method. This fits closely with the random walk hypothesis for the dividend–price ratio. Thus, the return forecast equals the sum of the current dividend–price ratio and the long-run historical average of earnings growth.[2]

We apply the SOP method using the same data as Goyal and Welch (2008) for the 1927–2007 period.[3] Our approach clearly performs better than both the historical mean and the traditional predictive regressions. We obtain an out-of-sample R-square (relative to the historical mean) of 1.32% with monthly data and 13.43% with yearly data (and nonoverlapping observations). This compares with out-of-sample R-squares ranging from −1.78% to 0.69% (monthly) and from −17.57% to 7.54% (yearly) obtained using the predictive regression approach in Goyal and Welch (2008).

The SOP method can be interpreted as a predictive regression with the dividend–price ratio as a predictor and with the restrictions that the intercept equals the historical average of earnings growth and the slope equals one. An important concern with our findings is that we might have picked, by chance, coefficients that are close to the in-sample estimates of the unrestricted predictive regression over the forecasting period. Then the out-of-sample R-square would really be an in-sample R-square. We address this concern by estimating the predictive regression and find that the in-sample estimated coefficients are very different from the SOP method implicit assumptions. This dissipates the concern that the SOP method is based on mining the coefficients. Using restricted versions of the predictive regression, we show that both the dividend–price ratio and the earnings growth components are equally responsible for the performance of the SOP method. We also find that the performance of the SOP method is robust to alternative estimates of the persistence of the dividend–price ratio and of the average earnings growth.

The gain in out-of-sample performance of the SOP method relative to predictive regressions is mainly due to the absence of estimation error that comes from a return forecast equal to the sum of the current dividend–price ratio and the long-run historical average of earnings growth; i.e., there are no parameters to estimate. A parallel exists in the exchange rate predictability literature. Meese and Rogoff (1983) and numerous authors since show that predictive regressions on fundamentals such as interest rate differentials cannot beat the random walk alternative out of sample. However, the literature on carry strategies shows that buying high interest rate and shorting low interest rate currencies produces consistent profits (see Burnside, Eichenbaum, and Rebelo, 2007; Burnside, Eichenbaum, Kleshchelski, and Rebelo, 2010) In a sense, these trading strategies predict exchange rates with interest rates but do not require any estimation and therefore have no estimation error.

Our results are robust in subsamples and in international data. The SOP method performs remarkably well on data from the UK and Japan, where predictability in stock returns is even stronger than in the US. The economic gains from a trading strategy that uses the simplest version of the SOP method are substantial. Its certainty equivalent gain is 1.8% per year, and the Sharpe ratio is more than 0.3% higher than a trading strategy based on the historical mean. In contrast, trading strategies based on predictive regressions would have generated significant economic losses. We conclude that there is substantial predictability in stock returns and that it would have been possible to profitably time the market in real time.

We conduct a Monte Carlo simulation experiment to better understand the performance of the SOP method.

[2] We also use two alternatives to predict the growth rate in the price–earnings ratio. In the first alternative, we use predictive regressions for the growth rate in the price–earnings ratio. In the second alternative, we regress the price–earnings ratio on macroeconomic variables and calculate the growth rate that would make the currently observed ratio revert to the fitted value. There is some improvement in the out-of-sample performance of the SOP method from using these alternatives.

[3] The sample period in Goyal and Welch (2008) is 1927–2004. We use the more recent data, but the results improve if we use only the 1927–2004 period.

We simulate the economy of Binsbergen and Koijen (2010), in which returns and dividend growth are assumed to be predictable. We find that the root mean square error (RMSE) of the simplest version of the SOP method (relative to the true expected return, which is known in the simulation) is 2.87%, compared with 4.94% for the historical mean and 3.73% for predictive regressions. The superior performance of the SOP estimator relative to the predictive regression estimator is explained by its lower variance due to the absence of estimation error. Relative to the historical mean, the SOP estimator presents similar variance, but much higher correlation with the true expected return.

The most important practical applications in finance—cost of capital calculation and portfolio management—require an estimate of stock market expected returns that works robustly out of sample with high explanatory power. Our paper offers the first estimator that meets these requirements. Going from out-of-sample *R*-squares that are close to zero in previous studies to *R*-squares of more than 13% matters hugely in practice.

2. Forecasting returns out of sample

We first describe the predictive regression methodology to forecast stock market returns. We then present a simple decomposition of stock returns and show how to forecast each component. Finally, we describe our main results.

2.1. Predictive regressions

The traditional predictive regression methodology regresses stock returns on lagged predictors[4]:

$$r_{t+1} = \alpha + \beta x_t + \varepsilon_{t+1}. \tag{1}$$

We generate out-of-sample forecasts of the stock market return using a sequence of expanding windows. Specifically, we take a subsample of the first s observations $t=1,\ldots,s$ of the entire sample of T observations and estimate regression Eq. (1). We denote the conditional expected return by $\mu_s = \mathrm{E}_s(r_{s+1})$, where $\mathrm{E}_s(\cdot)$ is the expectation operator conditional on the information available at time s. We then use the estimated coefficients of the predictive regression (denoted by hats) and the value of the predictive variable at time s to predict the return at time $s+1$[5]

$$\hat{\mu}_s = \hat{\alpha} + \hat{\beta} x_s. \tag{2}$$

We follow this process for $s=s_0,\ldots,T-1$, thereby generating a sequence of out-of-sample return forecasts $\hat{\mu}_s$. To start the procedure, we require an initial sample of size s_0 (20 years in the empirical application). This process simulates what a forecaster could have done in real time.

We evaluate the performance of the forecasting exercise with an out-of-sample *R*-square similar to the one proposed by Goyal and Welch (2008).[6] This measure compares the predictive ability of the regression with the historical sample mean:

$$R^2 = 1 - \frac{MSE_P}{MSE_M}, \tag{3}$$

where MSE_P is the mean square error of the out-of-sample predictions from the model:

$$MSE_P = \frac{1}{T-s_0} \sum_{s=s_0}^{T-1} (r_{s+1} - \hat{\mu}_s)^2, \tag{4}$$

and MSE_M is the mean square error of the historical sample mean:

$$MSE_M = \frac{1}{T-s_0} \sum_{s=s_0}^{T-1} (r_{s+1} - \bar{r}_s)^2, \tag{5}$$

where $\bar{r}_s$ is the historical mean of stock market returns up to time s.[7]

The out-of-sample *R*-square takes positive (negative) values when the model predicts returns better (worse) than the historical mean. Goyal and Welch (2008) offer evidence (replicated below) that predictive regressions using most variables proposed in the literature perform poorly out of sample.

We evaluate the statistical significance of the results using the *MSE–F* statistic proposed by McCracken (2007), which tests for the equality of the MSE of the unconditional (historical mean) and conditional (model) forecasts:

$$MSE-F = (T-s_0)\left(\frac{MSE_M - MSE_P}{MSE_P}\right). \tag{6}$$

The fitted value from a regression is a noisy estimate of the conditional expectation of the left-hand-side variable. This noise arises from the sampling error inherent in estimating model parameters using a finite (and often limited) sample. Because a regression tries to minimize squared errors, it tends to overfit in-sample. That is, the regression coefficients are calculated to minimize the sum of squared errors that arise both from the fundamental relation between the variables and from the sampling noise in the data. The second component is unlikely to hold robustly out of sample. Ashley (2006) shows that the unbiased forecast is no longer squared-error optimal in this setting. Instead, the minimum-MSE forecast represents a shrinkage of the unbiased forecast toward zero. This process squares nicely with a prior of no predictability in returns. We apply

[4] Alternatives to predictive regressions based on Bayesian methods, latent variables, analyst forecasts, and surveys have been suggested by Welch (2000), Claus and Thomas (2001), Brandt and Kang (2004), Pastor and Stambaugh (2009), and Binsbergen and Koijen (2010).

[5] To be more rigorous, we should index the estimated coefficients of the regression by s, $\hat{\alpha}_s$, and $\hat{\beta}_s$, as they change with the expanding sample. We suppress the subscript s for simplicity.

[6] See Diebold and Mariano (1995) and Clark and McCracken (2001) for alternative criteria to evaluate out-of-sample performance.

[7] Goyal and Welch (2008) include a degree-of-freedom adjustment in their *R*-square measure that we do not use. The purpose of adjusting a measure of goodness of fit for the degrees of freedom is to penalize in-sample overfit, which would likely worsen out-of-sample performance. Because the measure we use is already fully out of sample, there is no need for such adjustment. In any case, for the sample sizes and the number of explanatory variables used in this study, the degree-of-freedom adjustment would be minimal.

a simple shrinkage approach to the predictive regression coefficients in Eq. (2) suggested by Connor (1997), as described in Appendix A.[8]

2.2. Return components

We decompose the total return of the stock market index into dividend yield and capital gains:

$$1+R_{t+1} = 1+CG_{t+1}+DY_{t+1} = \frac{P_{t+1}}{P_t}+\frac{D_{t+1}}{P_t}, \tag{7}$$

where R_{t+1} is the return obtained from time t to time $t+1$, CG_{t+1} is the capital gain, DY_{t+1} is the dividend yield, P_{t+1} is the stock price at time $t+1$, and D_{t+1} is the dividend per share paid during the return period.[9]

The capital gains component can be written as

$$1+CG_{t+1} = \frac{P_{t+1}}{P_t} = \frac{P_{t+1}/E_{t+1}}{P_t/E_t}\frac{E_{t+1}}{E_t} = \frac{M_{t+1}}{M_t}\frac{E_{t+1}}{E_t} = (1+GM_{t+1})(1+GE_{t+1}), \tag{8}$$

where E_{t+1} denotes earnings per share at time $t+1$, M_{t+1} is the price–earnings multiple, GM_{t+1} is the price–earnings multiple growth rate, and GE_{t+1} is the earnings growth rate.

Instead of the price–earnings ratio, we could alternatively use any other price multiple such as the price–dividend ratio, the price-to-book ratio, or the price-to-sales ratio. In these alternatives, we must replace the growth in earnings by the growth rate of the denominator in the multiple (i.e., dividends, book value of equity, or sales).[10]

The dividend yield can in turn be decomposed as

$$DY_{t+1} = \frac{D_{t+1}}{P_t} = \frac{D_{t+1}}{P_{t+1}}\frac{P_{t+1}}{P_t} = DP_{t+1}(1+GM_{t+1})(1+GE_{t+1}), \tag{9}$$

where DP_{t+1} is the dividend–price ratio (which is distinct from the dividend yield in the timing of the dividend relative to the price).

Replacing the capital gain and the dividend yield in Eq. (7), we can write the total return as the product of the dividend–price ratio and the growth rates of the price–earnings ratio and earnings:

$$1+R_{t+1} = (1+GM_{t+1})(1+GE_{t+1})+DP_{t+1}(1+GM_{t+1})(1+GE_{t+1}) = (1+GM_{t+1})(1+GE_{t+1})(1+DP_{t+1}). \tag{10}$$

Finally, we make this expression additive by taking logs:

$$r_{t+1} = \log(1+R_{t+1}) = gm_{t+1}+ge_{t+1}+dp_{t+1}, \tag{11}$$

where lowercase variables denote log rates. Thus, log stock returns can be written as the sum of the growth in the price–earnings ratio, the growth in earnings, and the dividend–price ratio.

[8] Shrinkage has been widely used in finance for portfolio optimization problems but not for return forecasting. See Brandt (2009) for applications of shrinkage in portfolio management.

[9] Bogle (1991a, 1991b), Fama and French (1998), Arnott and Bernstein (2002), and Ibbotson and Chen (2003) offer similar decompositions of returns.

[10] In our empirical application we obtain similar findings using these three alternative price multiples.

2.3. The sum-of-the-parts method

We propose forecasting separately the components of the stock market return from Eq. (11):

$$\hat{\mu}_s = \hat{\mu}_s^{gm}+\hat{\mu}_s^{ge}+\hat{\mu}_s^{dp}. \tag{12}$$

We estimate the expected earnings growth $\hat{\mu}_s^{ge}$ using a 20-year moving average of the growth in earnings per share up to time s. This is consistent with the view that earnings growth is nearly unforecastable (Campbell and Shiller, 1988; Fama and French, 2002; and Cochrane, 2008) but has a low-frequency predictable component (Binsbergen and Koijen, 2010) possibly due to a change over time in inflation (earnings growth is a nominal variable).

The expected dividend–price ratio $\hat{\mu}_s^{dp}$ is estimated by the current dividend–price ratio dp_s (the logarithm of one plus the current dividend–price ratio). This implicitly assumes that the dividend–price ratio follows a random walk as Campbell (2008) proposes.

The choice of estimators for earnings growth and dividend–price ratio is not entirely uninformed. We could be criticized for choosing the estimators with knowledge of the persistence of earnings growth and dividend–price ratio. The concern is whether this would have been known to an investor in the beginning of the sample, for example in the 1950s. We therefore check in Section 2.5 the robustness of our results to using different window sizes for the moving average of earnings growth and estimating (out of sample) a first-order autoregression for the dividend–price ratio.

In the simplest version of the SOP method, we assume no multiple growth, i.e., $\hat{\mu}_s^{gm}=0$, which fits closely with the random walk hypothesis for the dividend–price ratio. The SOP method forecast at time s of the stock return at time $s+1$ can thus be written as

$$\hat{\mu}_s = \hat{\mu}_s^{ge}+\hat{\mu}_s^{dp}, \tag{13}$$

$$\hat{\mu}_s = \overline{g}_s+dp_s, \tag{14}$$

where $\overline{g}_s$ is the 20-year moving average of the growth in earnings per share up to time s and dp_s is the logarithm of one plus the current dividend–price ratio. This forecast looks like the traditional predictive regression:

$$r_{t+1} = \alpha+\beta dp_t+\varepsilon_{t+1}, \tag{15}$$

with the restrictions that the intercept α is set to $\overline{g}_s$ and the slope β is set to one.[11]

2.4. Results

We use the data set constructed by Goyal and Welch (2008), with monthly data to predict the monthly stock market return and yearly data (nonoverlapping) to predict the yearly stock market return.[12] The market return is proxied by the Standard & Poor's (S&P) 500 index

[11] We thank the referee for making this point and for suggesting the following analysis.

[12] Goyal and Welch (2008) forecast the equity premium, i.e., the stock market return minus the short-term riskless interest rate. In this paper, we forecast the market return but obtain similar results when we apply our approach to the equity premium.

Table 1
Summary statistics of return components.

This table reports mean, median, standard deviation, minimum, maximum, skewness, kurtosis, and first-order autocorrelation coefficient of the realized components of stocks market returns. *gm* is the growth in the price–earnings ratio. *ge* is the growth in earnings. *dp* is the dividend–price ratio. *r* is the stock market return. The sample period is from December 1927 through December 2007.

Panel A: Univariate statistics

Return components	Mean	Median	Standard deviation	Minimum	Maximum	Skewness	Kurtosis	AR(1)
Monthly frequency (December 1927–December 2007)								
gm	0.03	0.14	5.95	−30.41	36.71	0.05	9.74	0.16
ge	0.42	0.65	2.23	−9.52	15.12	−0.23	8.19	0.88
dp	0.33	0.31	0.14	0.09	1.27	1.15	6.84	0.98
r	0.79	1.26	5.55	−33.88	34.82	−0.43	11.19	0.08
Annual frequency (1927–2007)								
gm	0.44	−1.44	26.33	−62.26	78.83	0.27	3.12	−0.17
ge	5.09	9.64	21.49	−70.56	56.90	−1.02	5.42	0.17
dp	3.90	3.60	1.64	1.13	9.62	0.75	3.99	0.79
r	9.43	13.51	19.42	−60.97	43.60	−0.97	4.50	0.09

Panel B: Correlations

Return components	*gm*	*ge*	*dp*	*r*
Monthly frequency (December 1927–December 2007)				
gm	1			
ge	−0.35	1		
dp	−0.07	−0.20	1	
r	0.93	0.02	−0.13	1
Annual frequency (1927–2007)				
gm	1			
ge	−0.66	1		
dp	−0.21	−0.16	1	
r	0.60	0.19	−0.38	1

continuously compounded return including dividends. The sample period is from December 1927 to December 2007 (or 1927 to 2007 with annual data).

Table 1 presents summary statistics of stock market return (r) and its components (gm, ge, and dp) at the monthly and yearly frequency. The mean annual stock market return is 9.69% and the standard deviation is 19.42% over the whole sample period. Fig. 1 plots the yearly cumulative realized components of stock market return over time. Clearly average returns are driven mostly by earnings growth and the dividend–price ratio, while most of the return volatility comes from earnings growth and the price–earnings ratio growth. Fig. 1 shows that the time series properties of the return components are very different. The dividend–price ratio is very persistent, with an AR(1) coefficient of 0.79 at the annual frequency, while the AR(1) coefficients of earnings growth and multiple growth are close to zero.[13]

We perform an out-of-sample forecasting exercise along the lines of Goyal and Welch (2008). We examine the out-of-sample performance of a long list of predictors of stock returns. Appendix B provides a description of the predictors. Table 2 reports the results for the whole sample period. The forecast period starts 20 years after the beginning of the sample (January 1948) and ends in December 2007 for monthly frequency (1948–2007 for yearly frequency). Panel A reports results for monthly return forecasts, and Panel B reports results for annual return forecasts. Each row of the table uses a different forecasting variable. The asterisks in the in-sample R-square column denote significance of the in-sample regression as measured by the F-statistic. The asterisks in the out-of-sample R-squares columns denote whether the performance of the conditional forecast is statistically different from the unconditional forecasts (i.e., historical mean) using the McCracken (2007) *MSE–F* statistic.

The in-sample R-square of the full-sample regression in Panel A shows that most of the variables have modest predictive power for monthly stock returns over the long sample period considered here. The most successful variable is net equity expansion with an R-square of 1.07%. Overall, only four variables are significant at the 5% level.

The remaining two columns evaluate the out-of-sample performance of the different forecasts using the out-of-sample R-square relative to the historical mean. The fourth column reports the out-of-sample R-squares from the traditional predictive regression approach as in Goyal and Welch (2008). The fifth column reports the out-of-sample R-squares from the predictive regression with

[13] Earnings growth shows substantial persistence at the monthly frequency, but that is because we measure earnings over the previous 12 months, and there is therefore substantial overlap in the series from one month to the next.

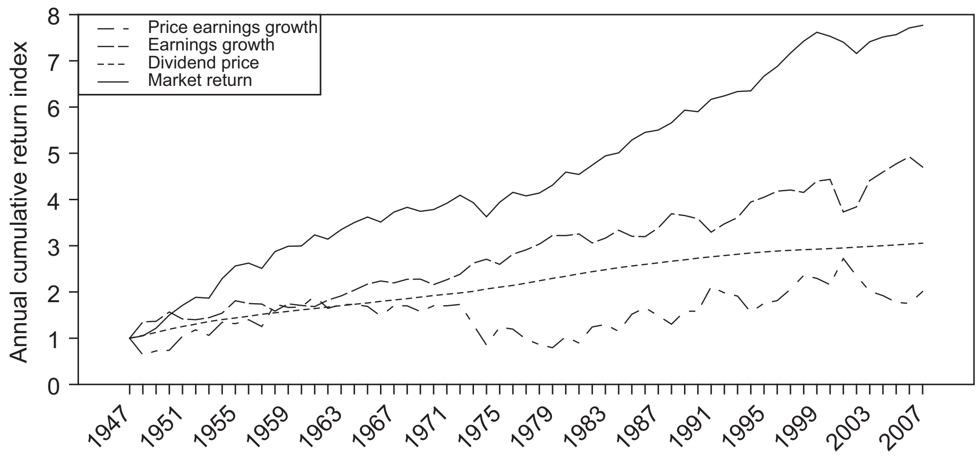

Fig. 1. Cumulative realized stock market components. This figure shows of annual realized price–earnings ratio growth (*gm*), earnings growth (*ge*), dividend price (*dp*), and stock market return (*r*) index (base year is 1947=1).

shrinkage. We present the out-of-sample *R*-squares from the sum-of-the-parts method with no multiple growth at the bottom of the panel.

Several conclusions stand out for the monthly return forecasts in Panel A. First, consistent with the findings in Goyal and Welch (2008), out-of-sample *R*-squares from the traditional predictive regression are in general negative, ranging from −1.78% to −0.05%. The one exception is the net equity expansion variable, which presents an out-of-sample *R*-square of 0.69% (significant at the 1% level).

Second, shrinkage improves the out-of-sample performance of most predictors. In the next column eight variables out of 16 have positive *R*-squares, although only two are significant at the 5% level. The *R*-squares are, however, still modest, with a maximum of 0.53%.

We next perform the out-of-sample forecasting exercise using the simplest version of the SOP method. Using only the dividend price and earnings growth components to forecast monthly stock market returns, we obtain an out-of-sample *R*-square of 1.32% (significant at the 1% level), which is much better than the performance of the traditional predictive regressions.

We now look at the annual stock market return forecasts. We use nonoverlapping returns to avoid the concerns with the measurement of *R*-squares with overlapping returns pointed out by Valkanov (2003) and Boudoukh, Richardson, and Whitelaw (2008). Our findings for monthly return forecasts are also valid at the annual frequency: Forecasting the components of stock market returns separately delivers out-of-sample *R*-squares significantly higher than traditional predictive regressions. There is an even more striking improvement at the yearly frequency.

The traditional predictive regression *R*-squares in Panel B are in general negative at yearly frequency (13 out of 16 variables)—consistent with Goyal and Welch (2008). The *R*-squares range from −17.57% to 7.54%, but only one variable is significant at the 1% level. Using shrinkage with traditional predictive regressions (next column) produces 11 variables with positive *R*-squares, but only two are significant at the 5% level. Forecasting the components of stock market returns separately dramatically improves performance. We obtain an *R*-square of 13.43% (significant at the 1% level) with the SOP method to forecast annual stock market returns in Panel B.

We compare the performance of the SOP method with the historical mean and predictive regression methods of forecasting stock market returns using graphical analysis. The aim is to understand better why the SOP method outperforms the alternative methods. We present and discuss the results at the annual frequency, but the conclusions are qualitatively similar using the monthly frequency.

Fig. 2 shows the SOP forecast of stock market return with no multiple growth and its two components. Substantial time variation exists in the stock market return forecasts over time, from nearly 4% per year around the year of 2000 to almost 15% per year in the early 1950s and the 1970s. The time variation of expected stock market return is due to both components. The SOP implicit equity premium in the middle panel also shows ample variability over time, ranging from approximately −2% to 13%. The equity premium was high in the early 1950s and slightly negative in the early 1980s when the stock market return forecast reached the highest figure but interest rates were also at record high levels. The bottom panel shows that the SOP forecast aligns with subsequent five-year average realized returns with the exception of the late 1940s (post-World War II economic growth surprise), early 1970s (oil shock), and mid-1990s (Internet bubble). In our forecast period, the average SOP return forecast of 9.52% is below the average realized stock market return of 11.28%, which is consistent with returns in this period having a positive surprise component. The SOP implicit equity premium is significantly negatively correlated with interest rates [Treasury bill rate (TBL) and term spread (TMS)] and positively correlated with the default yield spread (DFY). The *R*-square of a regression of the SOP equity premium on the default spread is 45% (untabulated results).

Fig. 3 compares the return forecasts from the SOP method with forecasts from traditional predictive regressions and the historical mean. There are large differences in the three forecasts. The expected returns using predictive regressions change drastically depending on the predictor. The historical mean and, to some extent, the predictive regressions tend to increase with past returns.

Table 2
Forecasts of stock market returns.

This table presents in-sample and out-of-sample *R*-squares (in percentage) for stock market return forecasts at monthly and annual (nonoverlapping) frequencies from predictive regressions and the sum-of-the-parts (SOP) method with no multiple growth. The in-sample *R*-squares are estimated over the full-sample period. The out-of-sample *R*-squares compare the forecast error of the model with the forecast error of the historical mean. The sample period is from December 1927 through December 2007. Forecasts begin 20 years after the sample start. Asterisks denote significance of the in-sample regression as measured by the *F*-statistic or significance of the out-of-sample *MSE–F* statistic of McCracken (2007). ***, **, and * denote significance at the 1%, 5%, and 10% level, respectively.

Variable	Description	In-sample *R*-square	Out-of-sample *R*-square	
Panel A: Monthly return forecasts (January 1948–December 2007)				
Predictive regressions			No shrinkage	Shrinkage
SVAR	Stock variance	0.05	−0.10	−0.02
DFR	Default return spread	0.08	−0.35	−0.05
LTY	Long-term bond yield	0.02	−1.19	−0.09
LTR	Long-term bond return	0.17	−0.98	−0.05
INFL	Inflation	0.04	−0.07	−0.02
TMS	Term spread	0.08	−0.05	0.04
TBL	T-bill rate	0.00	−0.59	−0.10
DFY	Default yield spread	0.03	−0.21	−0.03
NTIS	Net equity expansion	1.07***	0.69***	0.50**
ROE	Return on equity	0.07	−0.05	0.03
DE	Dividend payout	0.34*	−0.63	0.11
EP	Earnings price	0.76***	−0.51	0.53**
SEP	Smooth earning price	0.74**	−1.25	0.02
DP	Dividend price	0.15	−0.18	0.04
DY	Dividend yield	0.23	−0.58	0.07
BM	Book-to-market	0.58**	−1.78	−0.06
SOP with no multiple growth			1.32***	
Panel B: Annual return forecasts (1948–2007)				
Predictive regressions			No shrinkage	Shrinkage
SVAR	Stock variance	0.34	−0.15	0.00
DFR	Default return spread	1.95	1.64*	0.99
LTY	Long-term bond yield	0.71	−8.31	−0.85
LTR	Long-term bond return	2.29	−2.94	2.65**
INFL	Inflation	1.39	−1.04	0.53
TMS	Term spread	0.80	−7.23	−1.20
TBL	T-bill rate	0.13	−11.69	−2.09
DFY	Default yield spread	0.03	−1.13	−0.31
NTIS	Net equity expansion	12.29***	1.06*	2.30*
ROE	Return on equity	0.02	−10.79	−2.40
DE	Dividend payout	1.58	−0.17	0.47
EP	Earnings price	5.69**	7.54***	4.56**
SEP	Smooth earning price	8.27 **	−17.57	2.47*
DP	Dividend price	1.63	−1.01	0.28
DY	Dividend yield	2.31	−17.21	1.45*
BM	Book-to-market	5.76**	−8.80	0.82
SOP with no multiple growth			13.43***	

Thus, after a large run-up in the market (as in 1995–2000), these two methods forecast higher returns. The opposite happens with the SOP method because after a run-up the dividend–price ratio tends to be low. The first panel displays return forecasts from a predictive regression on the dividend–price ratio. It is interesting to notice the large difference between this forecast and the SOP forecast that relies on the same dividend–price ratio. This difference is driven by the regression coefficient estimates, which is analyzed in detail in Section 2.5.

Fig. 4 shows cumulative out-of-sample *R*-squares for both the SOP method and predictive regressions. The SOP method dominates over the whole sample period, with good fit, although a drop in predictability is evident over time.

2.5. Discussion

The SOP method looks like a predictive regression with the dividend–price ratio as a predictor and with the restrictions that the intercept equals the earnings growth historical average and the slope equals one. A concern with the SOP method is that we might have picked, by chance, coefficients that are close to the in-sample estimates of the predictive regression Eq. (15) over the forecasting period. In that case, the out-of-sample *R*-square would really be an in-sample *R*-square.

We address this concern by estimating the (in-sample) predictive regression Eq. (15) with annual returns over the forecasting period 1948–2007. We obtain an α estimate of -0.018, which compares with a $\overline{g}$ of 0.062 (average in the 1948–2007 period), and a β estimate of 3.747, which compares with a β of one in the SOP method. Clearly the coefficients of the predictive regression are very different from the coefficients implicit in the SOP method. To compare the explanatory power of this predictive regression with the out-of-sample results obtained with the SOP method, we compute a pseudo-out-of-sample *R*-square of the predictive regression relative to the historical mean. That is, we plug the fitted values of the predictive regression estimated with the full sample into Eq. (4). This can be interpreted as the maximum out-of-sample *R*-square that could be obtained if we knew the optimal parameters, which would be infeasible in real time. This pseudo-out-of-sample *R*-square for the predictive regression is 18.26%, which compares with an out-of-sample *R*-square of 13.43% for the SOP method.[14] This dissipates the concern that the SOP method is based on mining the coefficients but shows that despite a substantial difference from the optimal coefficients the out-of-sample explanatory power of the SOP method is close to the upper bound given by the (in-sample) predictive regression.

We next investigate which component, earnings growth or dividend–price ratio, is the main driver of the performance of the SOP method. We estimate the predictive regression Eq. (15) imposing alternatively the restrictions that $\alpha=\overline{g}=0.062$ or $\beta=1$. When we restrict the intercept α, we obtain a β estimate of 1.772 and the pseudo out-of-sample *R*-square is 14.91%. When we restrict the slope β, we obtain an α estimate of 0.078

[14] The true out-of-sample performance of the predictive regression estimated with an expanding sample is still worse than the historical mean. The out-of-sample *R*-square is -0.02%, which is slightly different from the number reported in Table 2 because we are using here the log of one plus the dividend–price ratio instead of the log of the dividend–price ratio used in Goyal and Welch (2008).

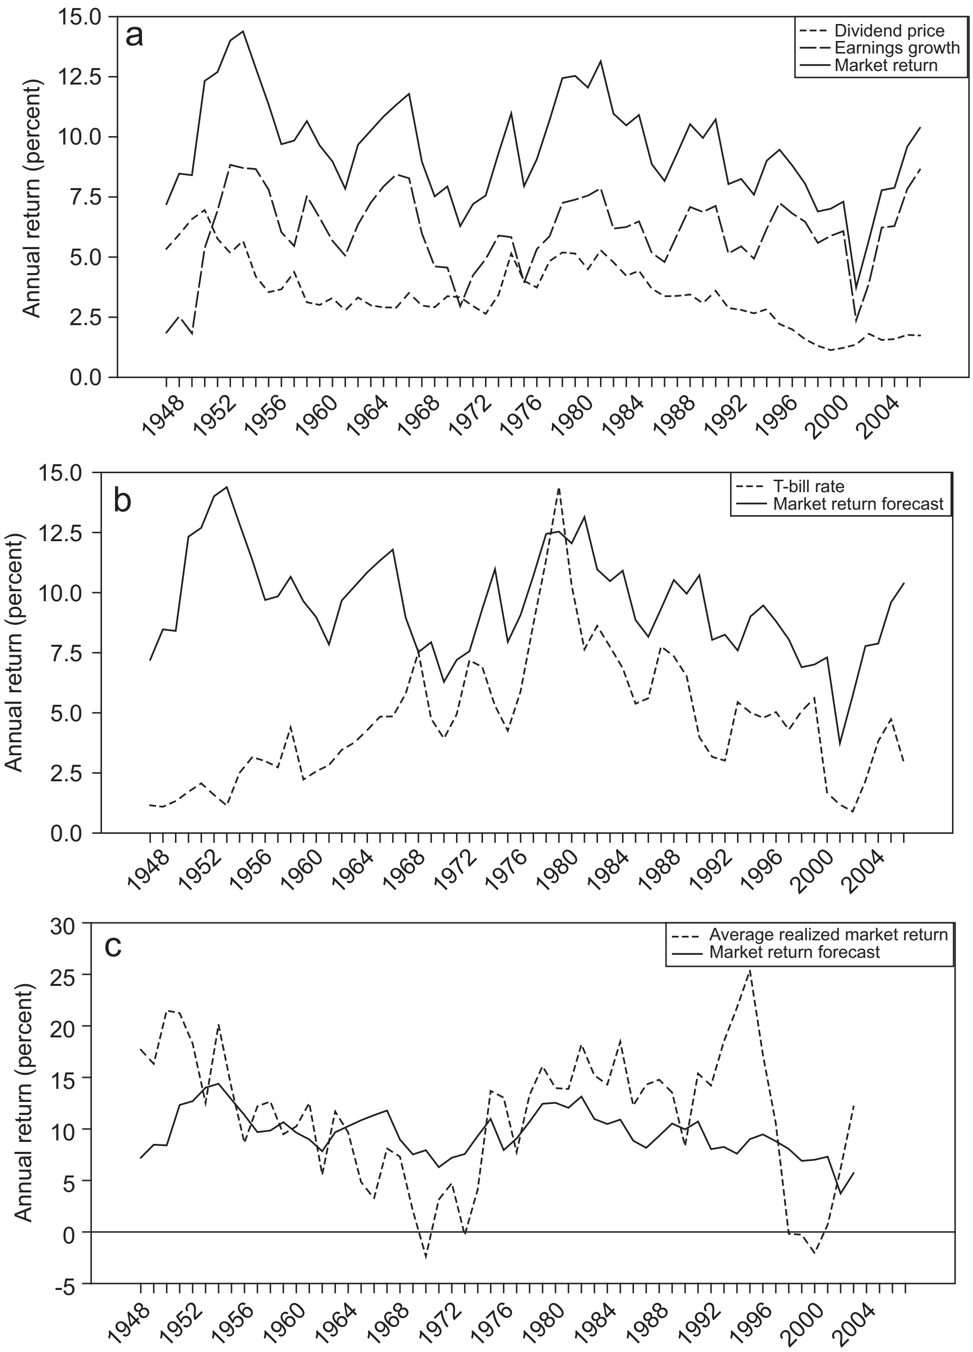

Fig. 2. SOP stock market return forecast. The top panel shows the forecast of earnings growth (*ge*), dividend price (*dp*), and market return (*ge*+*dp*) from the sum-of-the-parts (SOP) method with no multiple growth. The middle panel shows the Treasury bill rate (TBL) and the market return forecast from the SOP method with no multiple growth. The bottom panel shows the market return forecast from the SOP method with no multiple growth and the subsequent five-year average realized market return.

and the pseudo out-of-sample *R*-square is 12.65%.[15] The drop in *R*-square relative to the unrestricted predictive regression Eq. (15) is similar in both cases. We conclude that both components are responsible for the performance of the SOP method.

The same conclusion is supported by a variance decomposition of expected returns. We calculate the share of each component in the variance of expected returns estimated from the SOP method:

$$1 = \frac{\text{Var}(\overline{g}_s)}{\text{Var}(\hat{\mu}_s)} + \frac{\text{Var}(dp_s)}{\text{Var}(\hat{\mu}_s)} + \frac{2\,\text{Cov}(\overline{g}_s, dp_s)}{\text{Var}(\hat{\mu}_s)}. \tag{16}$$

The share of earnings growth is 56%, and the share of the dividend–price ratio is 43%. The covariance between the

[15] The true out-of-sample *R*-squares of the constrained regressions are 5.09% and 1.96%, respectively. There is therefore a gain relative to the unconstrained regression that reflects the lower estimation error in the constrained regressions. Still, the out-of-sample *R*-squares are substantially lower than the one obtained with the SOP method with no multiple growth, 13.43%.

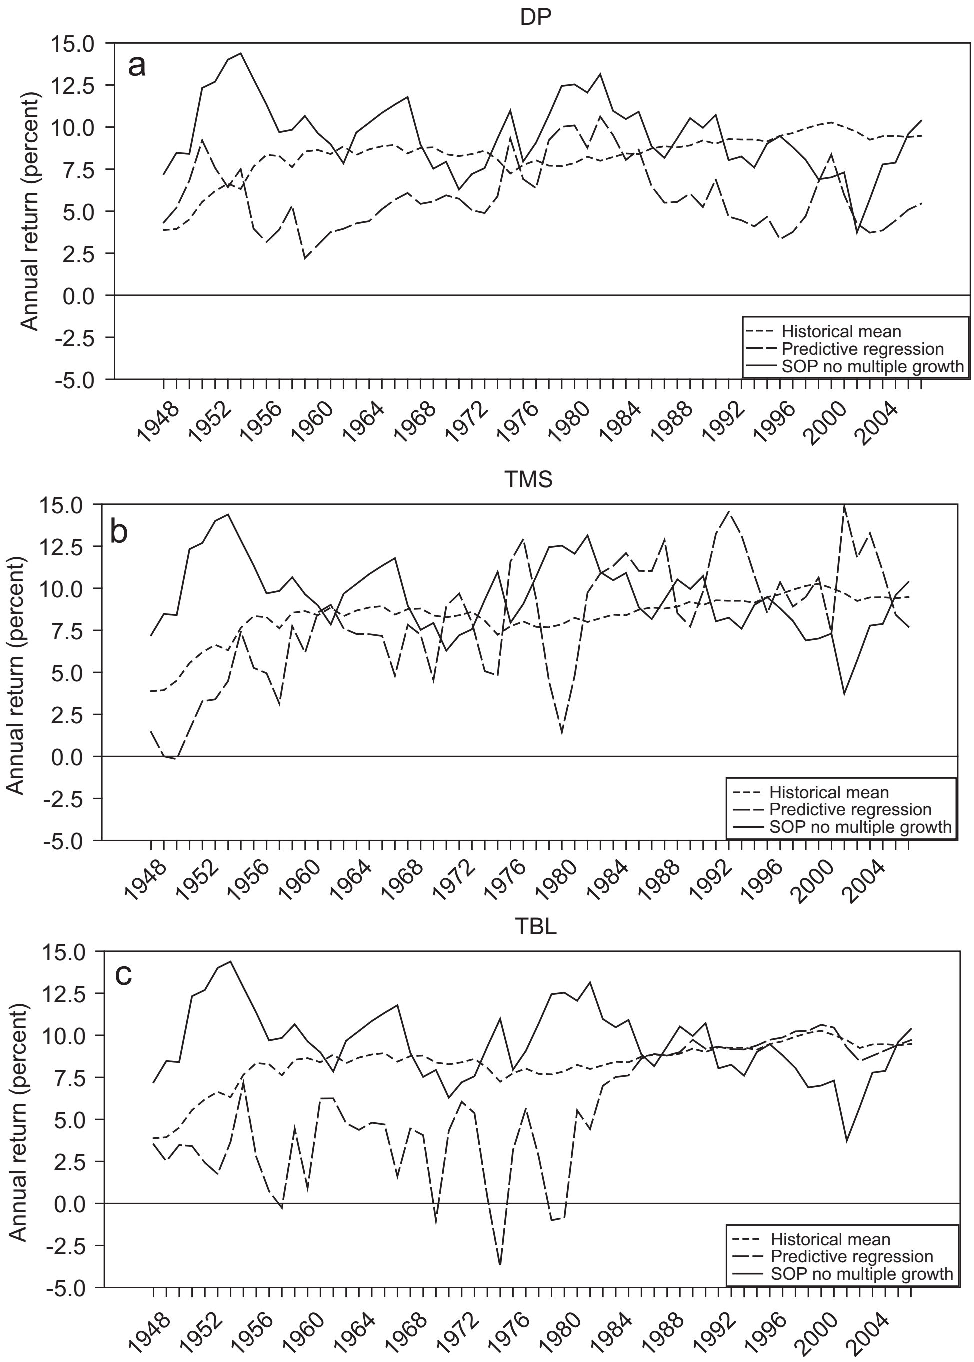

Fig. 3. Forecast of stock market return alternative methods. The panels show the forecast of market return from the historical mean, predictive regressions with dividend price (DP), term spread (TMS), and T-bill rate (TBL) as predictors, and sum-of-the-parts (SOP) method with no multiple growth.

earnings growth and the dividend–price ratio has a share of only 1%. These results indicate that both components contribute significantly and approximately with the same magnitude to the estimated expected returns. Our results are not driven by a single component.

The SOP method assumes that the persistence of the dividend–price ratio is very high and that the persistence of earnings growth is close to nil. This is implicit in forecasting the future dividend–price ratio with the current level of the ratio and in forecasting earnings growth with a 20-year moving average. Thus, another concern with the SOP method is that investors might not have known these facts about the persistence of earnings growth and of the dividend–price ratio at earlier times in the sample. In that case, the SOP method could not have been used in real time since the 1950s. We now investigate the sensitivity of the SOP method to these implicit assumptions.

We have used so far a 20-year moving average to estimate earnings growth. If we use alternatively ten- and 15-year moving averages, the resulting out-of-sample R-squares are virtually identical (12.20% and 13.46%, respectively). If we

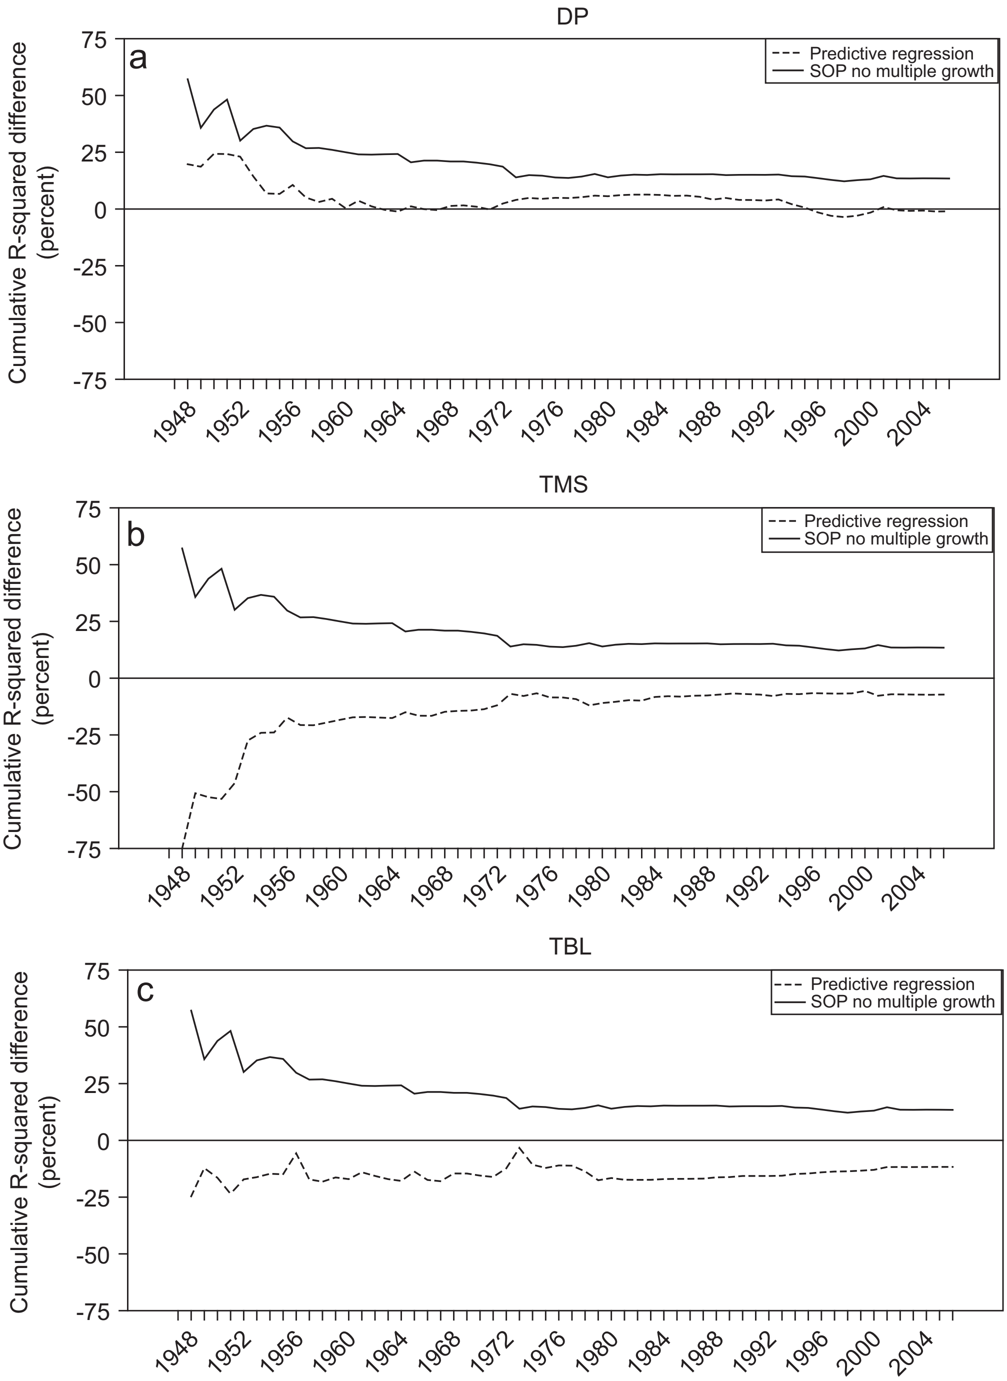

Fig. 4. Cumulative *R*-square. The panels show out-of-sample cumulative *R*-square up to each year from predictive regressions with dividend price (DP), term spread (TMS), and T-bill rate (TBL) as predictors, and the sum-of-the-parts (SOP) method with no multiple growth relative to the historical mean.

estimate earnings growth with an expanding window from the beginning of the sample, there is a slight deterioration of the *R*-square to 7.72%. This lower *R*-square indicates that there are low-frequency dynamics in earnings growth that are captured by the moving averages. We also try estimating a first-order autoregressive process for earnings growth, but we find that the coefficient is never significant and therefore such a specification would not have been chosen by investors at any time.

Instead of simply using the current level of the dividend–price ratio, we try estimating a first-order autoregressive process for the ratio and using the resulting forecast out of sample. We find that the autoregressive coefficient increases throughout the sample, from 0.4 in the beginning to 0.8 in the end, whereas the SOP method assumes implicitly that the coefficient is equal to one throughout. Substantially less persistence was evident in the dividend–price ratio earlier in the sample, and a legitimate concern exists that investors back then would not have modeled the dividend–price ratio as a random walk. However, when we use the autoregressive forecast of the dividend–price ratio as an alternative to the current ratio, we obtain an out-of-sample R-square as high as before, 12.47%. We conclude that the SOP method still works well even if we take into account a level of persistence in the dividend–price ratio substantially below a unit root.

It is instructive to compare our results with those in Campbell and Thompson (2008). They show that imposing restrictions on the signs of the coefficients of the predictive regressions modestly improves out-of-sample performance in both statistical and economic terms. More important, they suggest using the Gordon growth model to decompose expected stock returns (where earnings growth is entirely financed by retained earnings). Their method is a special case of our Eq. (12) with $\hat{\mu}_s^{gm}=0$ and $\hat{\mu}_s^{ge}=ROE_s(1-\mathrm{DE}_s)$, i.e., expected plowback times return on equity. The last component assumes that earnings growth corresponds to retained earnings times the return on equity. It is implicitly assumed that there are no external financing flows and that the marginal investment opportunities earn the same as the average return on equity.

Campbell and Thompson (2008) use historical averages to forecast the plowback (or one minus the payout ratio) and the return on equity. We implement their method in our sample, and the out-of-sample R-square is 0.54% (significant at the 5% level) with monthly frequency and 3.24% (significant only at the 10% level) with yearly frequency.[16] Our method using only the dividend–price ratio and earnings growth components gives significantly higher R-squares: 1.32% with monthly frequency and 13.43% with yearly frequency. In summary, the SOP substantially improves the out-of-sample explanatory power relative to previous studies, and the magnitude of this improvement is economically meaningful for investors.

3. Extensions and robustness

We use two alternative methods to forecast the growth in the price–earnings ratio. In the first approach, we run a traditional predictive regression—multiple growth regression—with the multiple growth gm (instead of the stock market return r) as the dependent variable:

$$gm_{t+1}=\alpha+\beta x_t+\varepsilon_{t+1}, \tag{17}$$

to obtain a forecast of the price–earnings ratio growth. We generate out-of-sample forecasts of the multiple growth using a sequence of expanding windows. As in the predictive regression approach, we apply shrinkage to the estimated coefficients as described in Appendix A.

The second approach—multiple reversion—assumes that the price–earnings ratio reverts to its expectation conditional on the state of the economy. We first run a time series regression of the multiple $m_t=\log M_t=\log(P_t/E_t)$ on the explanatory variable x_t:

$$m_t=a+bx_t+u_t. \tag{18}$$

[16] Campbell and Thompson (2008) use a longer sample period from 1891 to 2005 (with forecasts beginning in 1927) and obtain out-of-sample R-squares of 0.63% with monthly frequency and 4.35% with yearly frequency. We thank John Campbell for providing the data and programs for this comparison.

This is a contemporaneous regression as both sides of the equation are known at the same time. The fitted value of the regression gives the multiple that historically prevailed, on average, during economic periods characterized by the given level of the explanatory variable x. The expected value of the multiple at time s is

$$\hat{m}_s=\hat{a}+\hat{b}x_s. \tag{19}$$

If the observed multiple m_s is above this expectation, we anticipate negative growth for the multiple and vice versa. For example, suppose the current price–earnings ratio is 10 and the regression indicates that the expected value of the multiple is 12, given the current value of the explanatory variable. We would expect a return of 20% from this component. The estimated regression residual gives an estimate of the expected growth in the price multiple:

$$-\hat{u}_s=\hat{m}_s-m_s=\hat{\mu}_s^{gm}. \tag{20}$$

In practice, the reversion of the multiple to its expectation is quite slow and does not take place in a single period. To take this into account, we run a second regression of the realized multiple growth on the expected multiple growth using the estimated residuals from regression Eq. (18):

$$gm_{t+1}=c+d(-\hat{u}_t)+\nu_t. \tag{21}$$

Finally, we use these coefficients (after applying shrinkage as described in Appendix A) to forecast gm as

$$\hat{\mu}_s^{gm}=\hat{c}+\hat{d}(-\hat{u}_s). \tag{22}$$

We generate out-of-sample forecasts of the multiple growth using a sequence of expanding windows.

Table 3 reports the results for the whole sample period of these extensions of the SOP method. We use the same predictors to forecast the multiple growth gm in the sum-of-the-parts method than in predictive regressions. In the SOP method with multiple reversion we do not use the predictors that directly depend on the stock index price [earnings price (EP), smooth earning price (SEP), dividend price (DP), dividend yield (DY), and book-to-market (BM)] as this would correspond to explaining the price–earnings multiple with other multiples.

Panel A reports results for monthly return forecasts, and Panel B reports results for annual return forecasts. The R-squares in the SOP method with multiple growth regression in Panel A are all positive and range from 0.68% (book-to-market) to 1.55% (net equity expansion). Several variables turn in a good performance with R-squares above 1.3%, such as the term spread, inflation, T-bill rate, and the default yield spread. All the SOP method forecast results are significant at the 1% level under the McCracken (2007) *MSE–F* statistic.

The SOP method also presents good performance when we forecast the price–earnings growth using the multiple reversion approach. The last column shows that four (of 11 variables) have higher R-squares than in the SOP method with multiple growth regression. The R-square coefficients of the SOP method with multiple reversion range from 0.69% to 1.39%. The last figure in the last column gives the R-square of just using the historical mean of the price–earnings growth as a forecast of this

Table 3
Forecasts of stock market returns: sum-of-the-parts (SOP) extensions.

This table presents in-sample and out-of-sample *R*-squares (in percentage) for stock market return forecasts at monthly and annual (nonoverlapping) frequencies from the SOP method with multiple growth. The out-of-sample *R*-squares compare the forecast error of the model with the forecast error of the historical mean. The sample period is from December 1927 through December 2007. Forecasts begin 20 years after the sample start. Asterisks denote significance of the in-sample regression as measured by the *F*-statistic or significance of the out-of-sample *MSE–F* statistic of McCracken (2007). ***, **, and * denote significance at the 1%, 5%, and 10% level, respectively.

		Out-of-sample *R*-square	
Variable	Description	SOP with multiple growth regression	SOP with multiple reversion
Panel A: Monthly return forecasts (January 1948–December 2007)			
SVAR	Stock variance	0.91***	1.31***
DFR	Default return spread	1.27***	1.35***
LTY	Long-term bond yield	1.22***	0.69***
LTR	Long-term bond return	1.24***	1.35***
INFL	Inflation	1.37***	1.32***
TMS	Term spread	1.50***	1.39***
TBL	T-bill rate	1.31***	1.07***
DFY	Default yield spread	1.32***	1.34***
NTIS	Net equity expansion	1.55***	1.29***
ROE	Return on equity	1.20***	1.01***
DE	Dividend payout	1.20***	0.99***
EP	Earnings price	1.35***	–
SEP	Smooth earning price	0.94***	–
DP	Dividend price	0.89***	–
DY	Dividend yield	0.76***	–
BM	Book-to-market	0.68***	–
	Constant	–	1.35***
Panel B: Annual return forecasts (1948–2007)			
SVAR	Stock variance	12.74***	13.65***
DFR	Default return spread	14.40***	12.98***
LTY	Long-term bond yield	10.92***	7.61***
LTR	Long-term bond return	12.62***	16.94***
INFL	Inflation	12.91***	14.05***
TMS	Term spread	11.28***	15.57***
TBL	T-bill rate	11.51***	11.67***
DFY	Default yield spread	12.57***	14.46***
NTIS	Net equity expansion	13.31***	14.21***
ROE	Return on equity	13.66***	9.02***
DE	Dividend payout	12.60***	9.72***
EP	Earnings price	14.31***	–
SEP	Smooth earning price	11.07***	–
DP	Dividend price	8.99***	–
DY	Dividend yield	12.51***	–
BM	Book-to-market	10.20***	–
	Constant	–	14.40***

component, that is, assuming that the price earnings ratio reverts to its historical mean. We obtain a remarkable *R*-square of 1.35%.

Our findings for monthly return forecasts are also valid at the annual frequency when we use the extensions of the SOP method. When we add the forecast of the price–earnings growth from a predictive regression (SOP method with multiple growth regression) in Panel B, we obtain an even higher *R*-square for some variables: 14.31% (earnings price) and 14.40% (default return spread). And when we add the forecast of the price–earnings growth in the multiple reversion approach, the *R*-squares reach values of 16.94% (long-term bond return) and 15.57% (term spread). Under the SOP method, all variables are statistically significant at the 1% level. We conclude that, at yearly frequency, the SOP method with multiple reversion presents the best performance (compared with the SOP method with multiple growth regression) in a significant number of cases. This finding is not entirely surprising, as the speed of the multiple mean reversion is low.

Fig. 5 shows the realized price–earnings ratio and the fitted value from regressing the price–earnings on two different explanatory variables: the term spread (TMS) and the Treasury bill rate (TBL). This is one of the steps to obtain return forecasts in the SOP method with multiple reversion. It is interesting how little of the time variation of the price–earnings ratio is captured by these explanatory variables. It seems that the changes in the price–earnings ratio over time have little to do with the state of the economy. Important for our approach, we see that the realized price–earnings ratio reverts to the fitted value. This is not automatically guaranteed, because the forecasted price–earnings ratio is not the fitted value of a regression estimated ex post but is constructed from a series of regressions estimated with data up to each time. Yet the reversion is slow and at times takes almost ten years. The second regression in Eq. (21) captures this speed of adjustment. The expected return coming from the SOP method with multiple reversion varies substantially over time and takes both positive and negative values.

Fig. 6 shows the three versions of the SOP forecasts with two alternative predictive variables (TMS and TBL). The forecasts under the SOP method with no multiple growth are the same in the three panels. The three versions of the SOP method are highly correlated, but the SOP with multiple reversion displays more variability.

3.1. Subperiods

As Goyal and Welch (2008) find that predictive regressions perform particularly poor in the last decades, we repeat our out-of-sample performance analysis using two subsamples that divide the forecasting period in halves: from January 1948 through December 1976 and from January 1977 through December 2007. As before, forecasts begin 20 years after the subsample start. Table 4 presents the results. Panels A and C present results using monthly returns; and Panels B and D results using annual returns (nonoverlapping).

Like Goyal and Welch (2008), we also find better out-of-sample performance in the first subsample (which includes the Great Depression and World War II) than in the second subsample (which includes the oil shock of the 1970s and the Internet bubble at the end of the 20th century). The sum-of-the-parts method dominates the

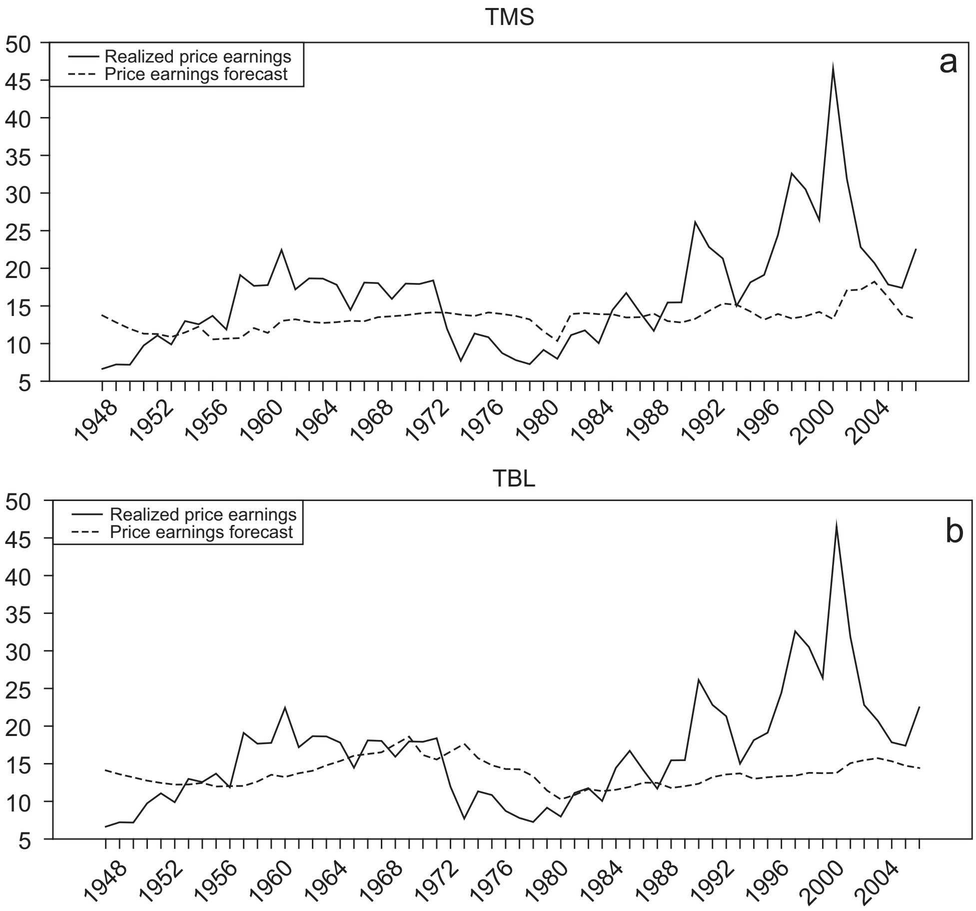

Fig. 5. Realized and forecasted price–earnings ratio. The panels show the realized price–earnings ratio and forecasted price–earnings ratio from the sum-of-the-parts (SOP) method with multiple reversion and using term spread (TMS) and T-bill rate (TBL) as predictors.

traditional predictive regressions in both subsamples and provides significant gains in performance over the historical mean.

Using monthly data, the out-of-sample *R*-squares of the traditional predictive regression are in general negative, ranging from -2.20% to 0.37% in the first subperiod and from -2.09% to 0.53% in the second subperiod. Net equity expansion has the best performance in both subperiods, and it is the only significant variable at the 5% level.

In both subperiods, a very significant improvement is evident in the out-of-sample forecasting performance when we separately model the components of the stock market return. As before, a considerable part of the improvement comes from the dividend–price ratio and earnings growth components alone: out-of-sample *R*-square of 1.80% in the first subperiod and 0.98% in the second subperiod (both significant at the 5% level) at the monthly frequency. The maximum *R*-squares using the SOP method with multiple growth regression are 2.29% in the first subperiod and 1.44% in the second subperiod (both significant at the 1% level). This is much better than the performance of the traditional predictive regressions. There is similar good performance when we use the SOP method with multiple reversion. The maximum *R*-squares are roughly 2% (in the first subperiod) and 1% (in the second subperiod), and they are all significant at the 5% level with one exception.

At the annual frequency, we find that most variables perform more poorly in the most recent subperiod. Using annual data, the out-of-sample *R*-squares of the traditional predictive regressions are in general negative in both subperiods. Forecasting the components of stock market returns separately, however, delivers positive and significant out-of-sample *R*-squares in both subperiods. As before, a considerable part of the improvement comes from the dividend–price ratio and earnings growth components alone. We obtain out-of-sample *R*-squares of 14.66% in the first subperiod and 12.10% in the second subperiod. The maximum *R*-squares using the SOP method with multiple growth regression are higher than 20% in the first subperiod and higher than 15% in the second subperiod (both significant at the 1% level). This is much better than the performance of the traditional predictive regressions.

3.2. *Trading strategies*

To assess the economic importance of the different approaches to forecast returns, we run out-of-sample trading strategies that combine the stock market with the risk-free asset. Each period, we use the different estimates of expected returns to calculate the Markowitz

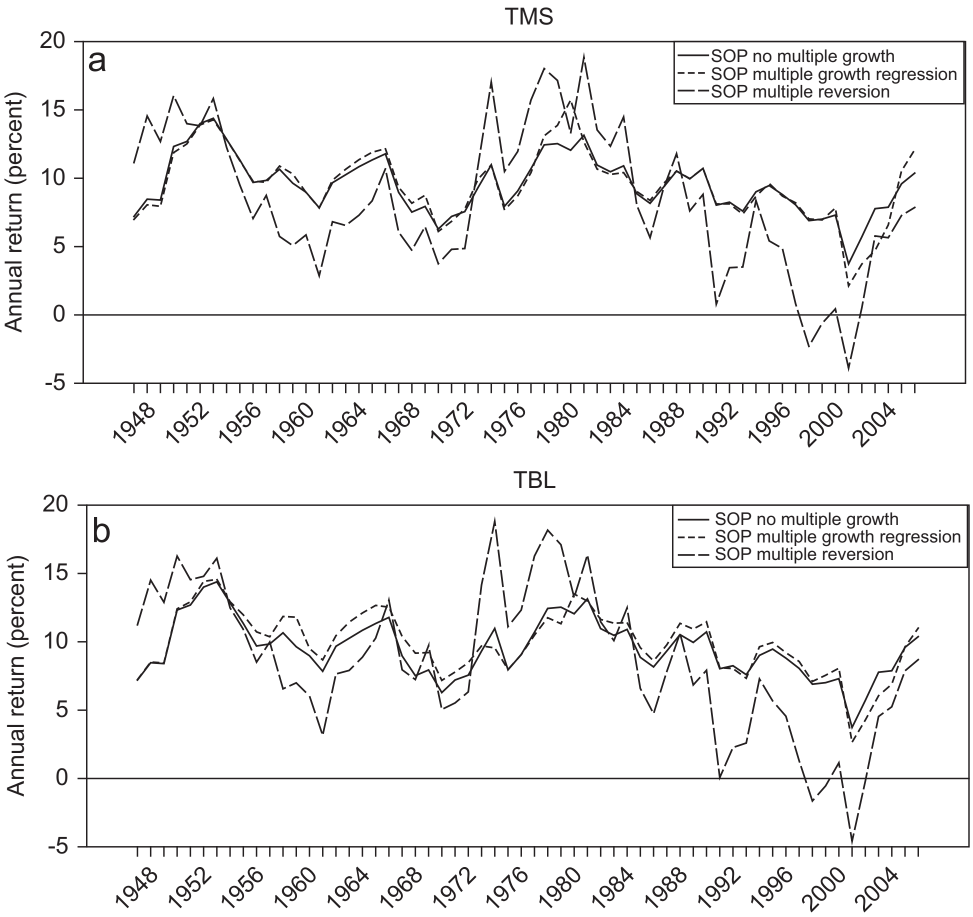

Fig. 6. Forecast of stock market return alternative sum-of-the-parts (SOP) methods. The panels show the forecast of market return from the SOP method with no multiple growth, and with multiple growth regression and multiple reversion using term spread (TMS) and T-bill rate (TBL) as predictors.

optimal weight on the stock market:

$$w_s = \frac{\hat{\mu}_s - rf_{s+1}}{\gamma \hat{\sigma}_s^2}, \tag{23}$$

where rf_{s+1} denotes the risk-free return from time s to $s+1$ (which is known at time s), γ is the risk-aversion coefficient assumed to be 2, and $\hat{\sigma}_s^2$ is the variance of the stock market returns that we estimate using all the available data up to time s.[17] The only thing that varies across portfolio policies are the estimates of the expected returns either from the predictive regressions or the sum-of-the-parts method. These portfolio policies could have been implemented in real time with data available at the time of the decision.[18]

We then calculate the portfolio return at the end of each period as

$$rp_{s+1} = w_s r_{s+1} + (1-w_s) rf_{s+1}. \tag{24}$$

We iterate this process until the end of the sample T, thereby obtaining a time series of returns for each trading strategy.

To evaluate the performance of the strategies, we calculate their certainty equivalent return:

$$ce = \overline{rp} - \frac{\gamma}{2}\sigma^2(rp), \tag{25}$$

where $\overline{rp}$ is the sample mean portfolio return and $\sigma^2(rp)$ is the sample variance portfolio return. This is the risk-free return that a mean–variance investor with a risk-aversion coefficient γ would consider equivalent to investing in the strategy. The certainty equivalent gain can also be interpreted as the fee the investor would be willing to pay to use the information in each forecast model. We also calculate the gain in Sharpe ratio (annualized) for each strategy.

Table 5 reports the certainty equivalent gains (annualized and in percentage) relative to investing based on the historical mean. Using the historical mean, the certainty equivalents are 7.4% and 6.4% per year at monthly and yearly frequencies, respectively. Using traditional predictive regressions leads to losses compared with the historical mean in most cases. Applying shrinkage to the traditional predictive regression slightly improves the performance of the trading strategies.

The SOP method always leads to economic gains relative to the historical mean. In fact, using only the

[17] Given the average stock market excess return and variance, a mean–variance investor with risk-aversion coefficient of 2 would allocate all wealth to the stock market. This is therefore consistent with equilibrium with this representative investor. Results are similar when we use other values for the risk-aversion coefficient.

[18] In untabulated results, we obtain slightly better certainty equivalents and Sharpe ratio gains if we impose portfolio constraints preventing investors from shorting stocks ($w_s \geq 0\%$) and assuming more than 50% leverage ($w_s \leq 150\%$).

Table 4
Forecasts of stock market returns: subsamples.

This table presents in-sample and out-of-sample *R*-squares (in percentage) for stock market return forecasts at monthly and annual (nonoverlapping) frequencies from predictive regressions and the sum-of-the-parts (SOP) method. The in-sample *R*-squares are estimated over the full-sample period. The out-of-sample *R*-squares compare the forecast error of the model with the forecast error of the historical mean. The sample period is from December 1927 through December 2007. Forecasts begin 20 years after the sample start. The subsamples divide the forecast period in half. Asterisks denote significance of the in-sample regression as measured by the *F*-statistic or significance of the out-of-sample *MSE–F* statistic of McCracken (2007). ***, **, and * denote significance at the 1%, 5%, and 10% level respectively.

			Out-of-sample *R*-square				
Variable	Description	In-sample *R*-square	Predictive regression	Predictive regression (shrinkage)	SOP no multiple growth	SOP multiple growth regression	SOP multiple reversion
Panel A: Monthly return forecasts (January 1948–December 1976)							
	–	–	–	–	1.80***	–	–
SVAR	Stock variance	0.00	−0.18	−0.04	–	1.64***	2.13***
DFR	Default return spread	0.01	−1.04	−0.22	–	1.57***	2.10***
LTY	Long-term bond yield	0.11	−1.72	0.04	–	1.61***	0.93***
LTR	Long-term bond return	0.12	−2.20	−0.34	–	1.42***	2.21***
INFL	Inflation	0.13	0.21*	0.08	–	2.19***	2.23***
TMS	Term spread	0.12	0.24*	0.10	–	2.06***	2.18***
TBL	T-bill rate	0.17	−0.15	0.09	–	1.90***	1.64***
DFY	Default yield spread	0.01	−0.53	−0.09	–	1.80***	2.11***
NTIS	Net equity expansion	1.08**	0.37*	0.16	–	1.85***	2.12***
ROE	Return on equity	0.01	−0.17	−0.03	–	1.74***	2.25***
DE	Dividend payout	0.47*	−1.09	0.02	–	1.73***	2.16***
EP	Earnings price	1.07**	−0.40	0.65**	–	2.15***	–
SEP	Smooth earnings price	1.83***	−1.45	0.11	–	2.06***	–
DP	Dividend price	0.24	0.29*	0.20*	–	2.26***	–
DY	Dividend yield	0.47*	−0.07	0.33*	–	2.29***	–
BM	Book-to-market	1.62***	0.04	0.39*	–	2.28***	–
	Constant	–	–	–	–	–	2.14***
Panel B: Annual return forecasts (1948–1976)							
	–	–	–	–	14.66***	–	–
SVAR	Stock variance	0.19	−0.76	−0.21	–	13.93***	21.54***
DFR	Default return spread	2.34	4.52*	1.66	–	14.82***	20.17***
LTY	Long-term bond yield	0.70	−10.95	−0.82	–	9.62**	13.40***
LTR	Long-term bond return	6.82*	9.64*	5.10**	–	13.44***	25.18***
INFL	Inflation	1.49	−0.99	0.72	–	13.77***	22.18***
TMS	Term spread	1.91	−6.66	-0.68	–	12.80***	21.85***
TBL	T-bill rate	1.59	−12.14	−1.43	–	11.66**	18.31***
DFY	Default yield spread	0.05	−1.64	−0.43	–	14.56***	21.98***
NTIS	Net equity expansion	14.91***	0.65	0.31	–	14.59***	21.75***
ROE	Return on equity	0.91	−12.62	−1.93	–	14.73***	22.82***
DE	Dividend payout	1.30	−0.23	−0.12	–	13.09***	21.52***
EP	Earnings price	6.74*	14.14***	4.71*	–	21.77***	–
SEP	Smooth earnings price	22.44***	−10.42	5.91**	–	23.21***	–
DP	Dividend price	2.93	4.48*	1.82	–	21.02***	–
DY	Dividend yield	5.28	−17.74	4.34*	–	18.16***	–
BM	Book-to-market	14.73***	8.30***	4.41*	–	19.87***	–
	Constant	–	–	–	–	–	21.76***
Panel C: Monthly return forecasts (January 1977–December 2007)							
	–	–	–	–	0.98**	–	–
SVAR	Stock variance	0.36	−0.99	−0.22	–	0.00	0.81**
DFR	Default return spread	0.14	−0.02	0.00	–	1.00**	0.87**
LTY	Long-term bond yield	0.05	−0.74	−0.11	–	0.93**	0.87**
LTR	Long-term bond return	0.74**	−0.67	0.19*	–	1.17***	0.85**
INFL	Inflation	0.03	−0.78	−0.13	–	0.88**	0.98**
TMS	Term spread	0.46*	−1.63	−0.01	–	1.10***	0.90**
TBL	T-bill rate	0.02	−2.09	−0.26	–	0.85**	1.09***
DFY	Default yield spread	1.02**	−0.14	0.25*	–	1.01**	0.60**
NTIS	Net equity expansion	0.85**	0.53**	0.58**	–	1.44***	0.79**
ROE	Return on equity	0.12	−0.88	−0.09	–	0.62**	0.80**
DE	Dividend payout	0.00	−1.07	−0.17	–	0.74**	−0.19
EP	Earnings price	0.61*	0.30*	0.19*	–	0.87**	–
SEP	Smooth earnings price	0.58*	−0.53	0.11	–	0.62**	–
DP	Dividend price	0.56*	−1.01	0.08	–	0.32*	–
DY	Dividend yield	0.61*	−1.31	0.08	–	0.23*	–
BM	Book-to-market	0.17	−0.73	−0.08	–	0.57**	–
Constant		–	–	–	–	–	0.86**

Table 4 (*continued*)

Variable	Description	In-sample R-square	Out-of-sample R-square				
			Predictive regression	Predictive regression (shrinkage)	SOP no multiple growth	SOP multiple growth regression	SOP multiple reversion
Panel D: Annual return forecast (1977–2007)							
	–	–	–	–	12.10***	–	–
SVAR	Stock variance	0.71	−25.88	−1.32	–	10.83**	7.28**
DFR	Default return spread	3.15	−5.65	−0.58	–	13.56***	11.51**
LTY	Long-term bond yield	2.37	−3.39	−0.09	–	11.55***	7.37**
LTR	Long-term bond return	3.26	−21.50	−0.15	–	12.35***	10.09**
INFL	Inflation	2.24	−10.39	−0.80	–	9.64**	11.64***
TMS	Term spread	1.27	−15.70	−2.04	–	9.92**	10.75**
TBL	T-bill rate	0.51	−17.57	−2.53	–	9.24**	8.81**
DFY	Default yield spread	5.71*	−14.77	−0.22	–	8.83**	9.38**
NTIS	Net equity expansion	2.53	1.53	3.30*	–	10.76**	8.08**
ROE	Return on equity	0.45	−9.68	0.43	–	15.32***	5.13**
DE	Dividend payout	0.14	−9.82	−1.79	–	11.35***	−7.56
EP	Earnings price	8.42**	1.82	3.49*	–	9.83**	–
SEP	Smooth earnings price	7.11*	−12.50	1.39	–	5.85**	–
DP	Dividend price	6.97*	−26.89	0.62	–	0.01	–
DY	Dividend yield	5.69*	−15.74	0.52	–	6.60**	–
BM	Book-to-market	3.04	−11.16	−0.50	–	6.33**	–
	Constant	–	–	–	–	–	9.74**

dividend–price ratio and earnings growth components, we obtain an economic gain of 1.79% per year. The greatest gains in the SOP method with multiple growth regression and multiple reversion are 2.33% and 1.72% per year, respectively. We obtain similar results using annual (nonoverlapping) returns.

Table 6 reports the gains in Sharpe ratio over investing using the historical mean. For the historical mean, the Sharpe ratios are 0.45 and 0.30 at the monthly and annual frequency, respectively. We find once again that using traditional predictive regressions leads to losses compared with the historical mean in most cases. Applying shrinkage to the traditional predictive regression slightly improves the performance of the trading strategies.

The SOP method always leads to Sharpe ratio gains relative to the historical mean. In fact, using only the dividend–price ratio and earnings growth components (SOP method with no multiple growth), we obtain a Sharpe ratio gain of 0.31. The maximum gains in the multiple growth regression and multiple reversion approaches are 0.33 and 0.24. We obtain similar Sharpe ratio gains using annual (nonoverlapping) returns.

3.3. *International evidence*

We repeat the analysis using international data. We obtain data on stock price indices and dividends from Global Financial Data (GFD) for the UK and Japan, which are the two largest stock markets in the world after the US. The sample period is from 1950 through 2007, which is shorter than the one in Tables 2 and 3 because of data availability. We report results using stock market returns in local currency at the annual frequency, but we obtain similar results using returns at the monthly frequency or returns in US dollars. We consider three macroeconomic variables (long-term yield, term spread, and Treasury bill rate, also obtained from GFD) and the dividend yield as predictors because these are the variables that are available for the longest sample period. We apply here the sum-of-the-parts method using price–dividend ratio as the price multiple instead of the price–earnings ratio because earnings for the UK and Japan are available only for a shorter period.

Table 7 presents the results for the international data. Panels A and B present the results for the UK and Japan, and Panel C presents the results for the US in the comparable sample period (1950–2007) and using the price–dividend ratio as multiple. The traditional predictive regression R-squares are in general negative, consistent with our previous findings. The R-squares range from −47.54% to 3.12%, and none is significant at the 5% level. Using shrinkage with the traditional prediction regression improves performance, and the R-square for the dividend yield is now significant at the 5% level in the UK and Japan (at only the 10% level in the US).

Forecasting the components of stock market returns separately dramatically improves performance. We obtain R-squares of 10.73% and 12.14% (both significant at the 1% level) in the UK and Japan, respectively, when we use only the dividend–price ratio and dividend growth components to forecast stock market returns (SOP method with no multiple growth). When we add the forecast of the price–dividend ratio growth from a predictive regression (SOP method with multiple growth regression), we obtain an even higher R-square for some variables: 13.28% in the UK using the dividend yield. When we alternatively add the forecast of the price–dividend ratio growth from the multiple reversion approach, the R-squares reach values of more than 11% in the UK and in Japan. Under the SOP method with multiple reversion, all variables are statistically

Table 5
Trading strategies: certainty equivalent gains.
This table presents out-of-sample portfolio choice results at monthly and annual (nonoverlapping) frequencies from predictive regressions and the sum-of-the-parts (SOP) method. The numbers are the certainty equivalent gains (in percentage) of a trading strategy timing the market with different return forecasts relative to timing the market with the historical mean return. The certainty equivalent return is $\overline{rp}-(\gamma/2)\hat{\sigma}^2(rp)$ with a risk-aversion coefficient of $\gamma=2$. All numbers are annualized (monthly certainty equivalent gains are multiplied by 12). The sample period is from December 1927 through December 2007. Forecasts begin 20 years after the sample start.

Variable	Description	Predictive regression	Predictive regression (shrinkage)	SOP no multiple growth	SOP multiple growth regression	SOP multiple reversion
Panel A: Monthly return forecasts (January 1948–December 2007)						
	–	–	–	1.79	–	–
SVAR	Stock variance	−0.04	0.00	–	0.97	1.61
DFR	Default return spread	−0.26	−0.04	–	1.75	1.72
LTY	Long-term bond yield	−1.56	−0.29	–	1.76	1.26
LTR	Long-term bond return	−0.25	0.10	–	1.92	1.68
INFL	Inflation	−0.07	−0.02	–	1.86	1.65
TMS	Term spread	0.41	0.18	–	2.13	1.72
TBL	T-bill rate	−0.86	−0.18	–	1.75	1.38
DFY	Default yield spread	−0.19	−0.05	–	1.53	1.65
NTIS	Net equity expansion	2.14	0.94	–	2.33	1.59
ROE	Return on equity	0.28	0.17	–	1.69	1.18
DE	Dividend payout	1.40	0.57	–	1.56	0.94
EP	Earnings price	0.20	0.35	–	1.69	–
SEP	Smooth earnings price	−1.15	−0.41	–	0.73	–
DP	Dividend price	−0.84	−0.26	–	0.62	–
DY	Dividend yield	−1.21	−0.33	–	0.45	–
BM	Book-to-market	−2.58	−0.52	–	0.49	–
	Constant	–	–	–	–	1.69
Panel B: Annual return forecasts (1948–2007)						
	–	–	–	1.82	–	–
SVAR	Stock variance	0.12	0.04	–	1.66	1.54
DFR	Default return spread	0.48	0.20	–	2.07	1.51
LTY	Long-term bond yield	−1.05	−0.19	–	1.75	0.92
LTR	Long-term bond return	1.48	0.66	–	1.88	1.95
INFL	Inflation	−0.08	0.08	–	1.73	1.47
TMS	Term spread	−0.58	−0.08	–	1.52	1.84
TBL	T-bill rate	−1.48	−0.31	–	1.69	1.25
DFY	Default yield spread	−0.01	−0.01	–	1.58	1.65
NTIS	Net equity expansion	1.25	0.54	–	1.89	1.64
ROE	Return on equity	−1.09	−0.28	–	2.04	0.78
DE	Dividend payout	0.60	0.24	–	1.91	0.74
EP	Earnings price	0.58	0.34	–	1.66	–
SEP	Smooth earnings price	−1.39	−0.14	–	0.88	–
DP	Dividend price	−0.71	−0.22	–	0.54	–
DY	Dividend yield	−2.04	−0.16	–	1.41	–
BM	Book-to-market	−1.53	−0.27	–	0.97	–
	Constant	–	–	–	–	1.67

significant at the 5% level. The SOP method performs better in the UK and in Japan than in the US when we redo the analysis for the comparable sample period and using the price–dividend ratio as the multiple (Panel C). The SOP method clearly dominates predictive regressions in US data.

Fig. 7 shows forecasts of stock market return for the UK, Japan, and the US according to the SOP method (with no multiple growth). There are substantial differences. The UK generally offers the highest expected returns (around 11.8% on average), while expected returns in Japan are the lowest through most of the sample (4.7% on average). At times, the difference in return forecasts across countries is as high as 12 percentage points. There is more variability in return forecasts in the UK and Japan than in the US. The correlation between expected returns in the UK and the US is high (on the order of 0.7), but Japanese expected returns are negatively correlated with both the UK and US stock markets (on the order of −0.3).

3.4. Analyst forecasts

An alternative forecast of earnings can be obtained from analyst estimates drawn from Institutional Brokers' Estimate System (I/B/E/S) and aggregated across all S&P 500 stocks. We use these forecasts to calculate both the price–earnings ratio and the earnings growth. Panel A of Table 8 reports the results for the sample period from January 1982 (when I/B/E/S data start) through December 2007 with monthly frequency. In this exercise we begin forecasts five years after the sample start, instead of 20 years as we did before, because of the shorter sample.

Table 6
Trading strategies: Sharpe ratio gains.

This table presents out-of-sample portfolio choice results at monthly and annual (nonoverlapping) frequencies from predictive regressions and the sum-of-the-parts (SOP) method. The numbers are the change in Sharpe ratio of a trading strategy timing the market with different return forecasts relative to timing the market with the historical mean return. All numbers are annualized. The sample period is from December 1927 through December 2007. Forecasts begin 20 years after the sample start.

Variable	Description	Predictive regression	Predictive regression (shrinkage)	SOP no multiple growth	SOP multiple growth regression	SOP multiple reversion
Panel A: Monthly return forecasts (January 1948–December 2007)						
	–	–	–	0.31	–	–
SVAR	Stock variance	0.00	0.00	–	0.12	0.22
DFR	Default return spread	−0.06	−0.01	–	0.30	0.24
LTY	Long-term bond yield	−0.25	−0.06	–	0.29	0.09
LTR	Long-term bond return	−0.12	−0.02	–	0.23	0.24
INFL	Inflation	−0.04	−0.01	–	0.31	0.19
TMS	Term spread	−0.05	−0.02	–	0.28	0.23
TBL	T-bill rate	−0.18	−0.04	–	0.32	0.16
DFY	Default yield spread	−0.02	0.00	–	0.33	0.24
NTIS	Net equity expansion	0.04	0.06	–	0.28	0.24
ROE	Return on equity	−0.06	−0.02	–	0.27	0.12
DE	Dividend payout	−0.02	0.00	–	0.32	0.14
EP	Earnings price	−0.09	0.30	–	0.23	–
SEP	Smooth earnings price	−0.21	0.12	–	0.12	–
DP	Dividend price	0.11	0.08	–	0.14	–
DY	Dividend yield	−0.13	0.15	–	0.07	–
BM	Book-to-market	−0.34	0.04	–	0.01	–
	Constant	–	–	–	–	0.24
Panel B: Annual return forecasts (1948–2007)						
	–	–	–	0.22	–	–
SVAR	Stock variance	0.01	0.00	–	0.23	0.11
DFR	Default return spread	0.03	0.02	–	0.23	0.12
LTY	Long-term bond yield	−0.14	−0.03	–	0.19	0.02
LTR	Long-term bond return	0.08	0.06	–	0.23	0.15
INFL	Inflation	0.01	0.02	–	0.21	0.09
TMS	Term spread	−0.10	−0.02	–	0.18	0.15
TBL	T-bill rate	−0.19	−0.04	–	0.19	0.08
DFY	Default yield spread	−0.01	−0.01	–	0.24	0.13
NTIS	Net equity expansion	0.05	0.04	–	0.22	0.12
ROE	Return on equity	−0.15	−0.04	–	0.16	0.03
DE	Dividend payout	0.00	0.00	–	0.21	0.04
EP	Earnings price	0.05	0.15	–	0.12	–
SEP	Smooth earnings price	−0.15	0.07	–	0.06	–
DP	Dividend price	−0.02	0.02	–	0.04	–
DY	Dividend yield	−0.21	0.07	–	0.20	–
BM	Book-to-market	−0.19	0.03	–	0.09	–
	Constant	–	–	–	–	0.13

Panel B replicates the analysis of Tables 2 and 3 for the same sample period for comparison.

We find that analyst forecasts work well with out-of-sample R-squares between 1.70% and 3.10%. However, the SOP method based only on historical data growth works even better than based on analyst forecasts in this sample period, with out-of-sample R-squares between 2.81% and 4.66%. This is consistent with the well-known bias in analyst forecasts.

4. Simulation analysis

In this section, we conduct a Monte Carlo simulation experiment to better understand the performance of the sum-of-the-parts method. We simulate the economy in Binsbergen and Koijen (2010), in which expected returns (μ_t) and expected dividend growth rates (g_t) follow AR(1) processes:

$$\mu_{t+1} = \delta_0 + \delta_1(\mu_t - \delta_0) + \varepsilon^{\mu}_{t+1} \tag{26}$$

and

$$g_{t+1} = \gamma_0 + \gamma_1(g_t - \gamma_0) + \varepsilon^{g}_{t+1}. \tag{27}$$

The dividend growth rate is equal to the expected dividend growth rate plus an orthogonal shock:

$$\Delta d_{t+1} = g_t + \varepsilon^{d}_{t+1}. \tag{28}$$

The Campbell and Shiller (1988) log-linear present value model implies that

$$r_{t+1} = \kappa + \rho(p_{t+1} - d_{t+1}) + \Delta d_{t+1} - (p_t - d_t), \tag{29}$$

where κ and ρ are constants of the log linearization.

Table 7
Forecasts of stock market returns: international evidence.
This table presents in-sample and out-of-sample *R*-squares (in percentage) for stock market return forecasts at annual (nonoverlapping) frequency in the UK (Panel A), Japan (Panel B), and the US (Panel C) from predictive regressions and the sum-of-the-parts (SOP) method. The in-sample *R*-squares are estimated over the full-sample period. The out-of-sample *R*-squares compare the forecast error of the model with the forecast error of the historical mean. The sample period is from 1950 or 1960 through 2007. Forecasts begin 20 years after the sample start. Asterisks denote significance of the in-sample regression as measured by the *F*-statistic or significance of the out-of-sample *MSE–F* statistic of McCracken (2007). ***, **, and * denote significance at the 1%, 5%, and 10% level respectively.

				Out-of-sample *R*-square				
Variable	Description	Forecast start	In-sample *R*-square	Predictive regression	Predictive regression (shrinkage)	SOP no multiple growth	SOP multiple growth regression	SOP multiple reversion
Panel A: UK annual return forecasts								
	–	1970	–	–	–	10.73***	–	–
LTY	Long-term bond yield	1970	5.29*	−47.54	−5.61	–	4.16**	11.27***
TMS	Term spread	1970	3.10	−14.71	−1.13	–	9.26**	11.60***
TBL	T-bill rate	1970	1.47	−20.87	−3.07	–	6.39**	11.51***
DY	Dividend yield	1970	11.97***	−9.19	5.07**	–	13.28***	10.78***
	Constant	1970	–	–	–	–	–	11.75***
Panel B: Japan annual return forecasts								
	–	1970	–	–	–	12.14***	–	–
LTY	Long-term bond yield	1970	1.69	−11.01	−1.86	–	12.11***	11.87***
TMS	Term spread	1980	0.36	−5.46	−0.89	–	5.75**	5.82**
TBL	T-bill rate	1980	1.76	−7.57	−0.62	–	5.14*	5.62**
DY	Dividend yield	1970	15.24***	3.12*	6.63**	–	10.25***	11.99***
	Constant	1970	–	–	–	–	–	11.91***
Panel C: US annual return forecasts								
	–	1970	–	–	–	7.75**	–	–
LTY	Long-term bond yield	1970	0.17	−20.73	−1.51	–	4.47**	3.12*
TMS	Term spread	1970	1.11	−12.05	−0.99	–	8.24**	5.50**
TBL	T-bill rate	1970	0.03	−21.18	−2.00	–	5.06**	3.40*
DY	Dividend yield	1970	7.95**	0.96	2.68*	–	6.64**	5.73**
	Constant	1970	–	–	–	–	–	5.92**

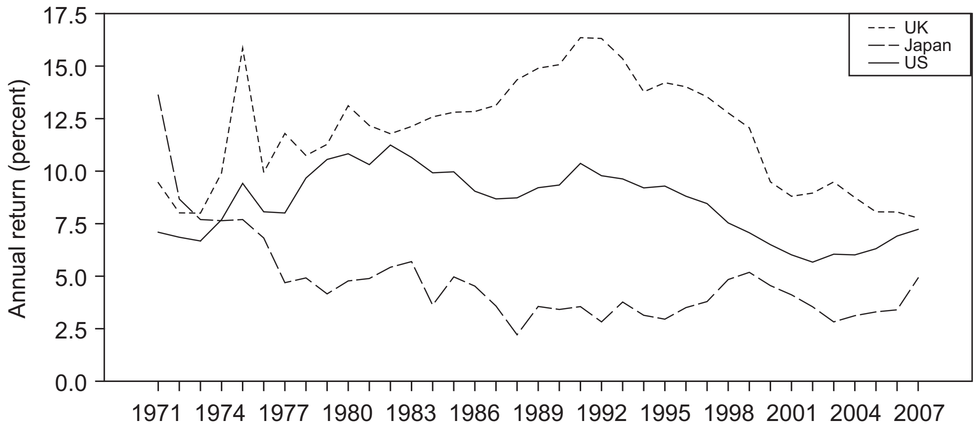

Fig. 7. Sum-of-the-parts (SOP) method forecast of international stock market returns. The figure shows annual forecast of market return in the UK, Japan, and the US from the SOP method with no multiple growth.

Iterating Eq. (29) and using processes Eqs. (26)– (28), it follows that

$$p_t - d_t = A - B_1(\mu_t - \delta_0) + B_2(g_t - \gamma_0), \tag{30}$$

where $A = \kappa/(1-\rho) + (\gamma_0 - \delta_0)/(1-\rho)$, $B_1 = 1/(1-\rho\delta_1)$, and $B_2 = 1/(1-\rho\gamma_1)$.

We simulate returns, dividend growth, and the dividend–price ratio from these equations using the estimated parameter values in Table II of Binsbergen and Koijen (2010). Because the AR(1) process for expected dividend growth rates in Eq. (27) can be written as an infinite moving average model, there is predictability of dividend growth by a smoothed average of past growth rates in this model.

We simulate ten thousand samples of 80 years (which is approximately the size of our empirical sample) of returns, dividend growth, and the dividend–price ratio for this economy. We use the simulated data to study the different forecasting methods of stock market return. The advantage of using Monte Carlo simulation is that we know the true expected return at each particular time. Thus, we can compare our forecasts with (true) expected returns and not just with realized returns as we do in the empirical analysis.

Table 8
Forecasts of stock market returns: analyst earnings forecasts

This table presents in-sample and out-of-sample R-squares (in percentage) for stock market return forecasts at monthly frequency from predictive regressions and the sum-of-the-parts (SOP) method. The in-sample R-squares are estimated over the full-sample period. The out-of-sample R-squares compare the forecast error of the model with the forecast error of the historical mean. The SOP method uses alternatively analyst earnings forecasts and historical earnings to calculate gm and ge. The sample period is from January 1982 through December 2007. Forecasts begin five years after the sample start. Asterisks denote significance of the in-sample regression as measured by the F-statistic or significance of the out-of-sample MSE–F statistic of McCracken (2007). ***, **, and * denote significance at the 1%, 5%, and 10% level respectively.

			Out-of-sample R-square							
					SOP—Analyst forecasts			SOP—Historical data		
Variable	Description	In-sample R-square	Predictive regression	Predictive regression (shrinkage)	No multiple growth	Multiple growth regression	Multiple reversion	No multiple growth	Multiple growth regression	Multiple reversion
Monthly return forecasts (January 1987–December 2007)										
	–	–	–	–	2.32***	–	–	3.62***	–	–
SVAR	Stock variance	0.88	−2.97	−0.17	–	2.26***	2.15***	–	2.81***	3.59***
DFR	Default return spread	0.60	−2.20	−0.14	–	2.20***	2.19***	–	3.51***	3.62***
LTY	Long-term bond yield	0.26	−0.67	−0.02	–	2.25***	3.10***	–	3.52***	4.66***
LTR	Long-term bond return	0.26	−0.45	−0.01	–	2.26***	2.17***	–	3.60***	3.58***
INFL	Inflation	0.04	−0.76	−0.06	–	2.22***	2.13***	–	3.54***	3.62***
TMS	Term spread	0.01	−2.00	−0.15	–	2.15***	2.07***	–	3.22***	3.59***
TBL	T-bill rate	0.29	−1.18	−0.03	–	2.19***	2.60***	–	3.37***	4.23***
DFY	Default yield spread	0.51	−0.49	0.04	–	2.12***	2.31***	–	3.34***	3.50***
NTIS	Net equity expansion	0.66	−1.23	0.06	–	2.01***	2.37***	–	2.97***	3.62***
ROE	Return on equity	0.02	−1.84	−0.10	–	1.97***	2.23***	–	3.17***	3.54***
DE	Dividend payout	0.02	−1.79	−0.12	–	2.26***	1.70***	–	3.59***	3.13***
EP	Earnings price	2.68***	1.78***	0.56*	–	2.39***	–	–	3.61***	–
SEP	Smooth earnings price	1.25**	−0.22	0.19	–	2.13***	–	–	3.39***	–
DP	Dividend price	1.74**	0.00	0.29*	–	2.08***	–	–	3.29***	–
DY	Dividend yield	1.74**	−0.23	0.28*	–	2.07***	–	–	3.25***	–
BM	Book-to-market	1.02*	0.14	0.14	–	2.16***	–	–	3.39***	–
	Constant	–	–	–	–	–	2.17***	–	–	3.61***

In each simulation of the economy, we replicate our out-of-sample empirical analysis; that is, we compute for each year the forecast of returns from the three approaches (historical mean, predictive regression, and SOP method with no multiple growth) using only past data. The regressions use the log dividend–price ratio as predictive variable.

Fig. 8 shows a scatter plot of each estimator of expected returns versus the true expected returns at the end of the simulated samples. The SOP expected return estimates have the lowest bias and variance.[19] The historical mean also has a low variance, but it does not capture the variation in the true expected returns. Predictive regressions have poor performance in terms of predicting stock market returns, with a higher variance than the SOP and historical mean methods.

[19] There is a slight "smile" shape in the relation between the true expected returns and the SOP estimates of expected returns. This happens because we simulate expected returns from a Campbell-Shiller approximation in which a linear relation exists between expected returns and the log of the dividend–price ratio. The corresponding relation in the decomposition underlying the SOP is between expected returns and the log of one plus the dividend price ratio. The difference between the log of the ratio and the log of one plus the ratio explains the nonlinearity in the plot.

To quantify this analysis, we compute the sum of the squares of the difference between the estimates of expected returns and the true expected returns from the simulation. Panel A of Table 9 presents the mean and the percentiles (across simulations) of the root mean square error of each forecast method. The results clearly show that the SOP method yields a better estimate of expected returns than both predictive regressions and the historical mean of returns. The average RMSE of the SOP method is 2.87%, which is low in absolute terms and significantly lower than the RMSE of the historical mean and predictive regressions, 4.94% and 3.73%, respectively. This difference persists across all the percentiles of the distribution of the RMSE.

We can decompose the expected MSE (across simulations) of each estimator of expected returns in the following way:

$$\frac{1}{T-s_0}\sum_{s=s_0}^{T-1} \mathrm{E}[(\hat{\mu}_s-\mu_s)^2] = \frac{1}{T-s_0}\sum_{s=s_0}^{T-1}[\mathrm{E}(\hat{\mu}_s-\mu_s)]^2+\frac{1}{T-s_0}\sum_{s=s_0}^{T-1}\mathrm{Var}(\hat{\mu}_s) + \frac{1}{T-s_0}\sum_{s=s_0}^{T-1}\mathrm{Var}(\mu_s)-\frac{1}{T-s_0}\sum_{s=s_0}^{T-1} 2\,\mathrm{Cov}(\hat{\mu}_s,\mu_s), \quad (31)$$

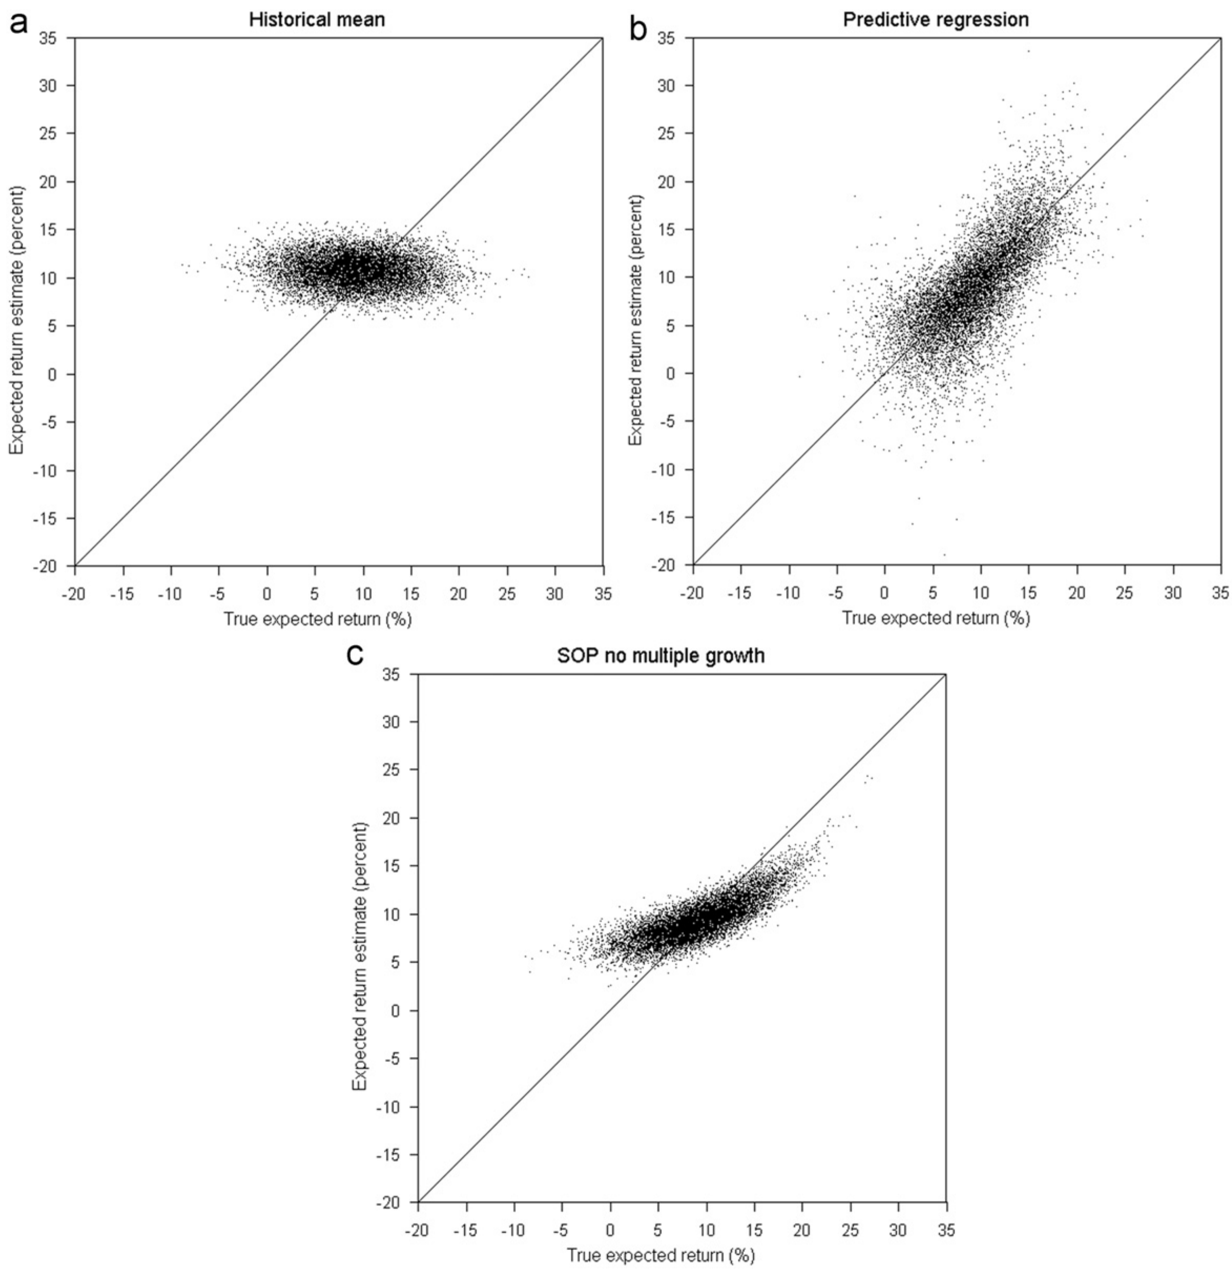

Fig. 8. Monte Carlo simulation. The expected return estimates are plotted against true expected returns from a Monte Carlo simulation of the economy in Binsbergen and Koijen (2010). The simulation generates ten thousand samples of 80 years of returns, dividend growth, and the dividend–price ratio for this economy. In each simulation of the economy, annual expected returns are estimated with past data using the historical mean, predictive regression with the log dividend–price ratio as conditioning variable, and the sum-of-the-parts (SOP) method with no multiple growth. The solid line is a 45° line.

where E(·), Var(·), and Cov(·) are moments across simulations. The first term corresponds to the square of the bias of the estimates of expected returns. The second term is the variance of the estimates of expected returns. The third term is the variance of the true expected return (which is the same for all methods). The final term is the covariance between the estimates of expected returns and the true expected return. Panel B of Table 9 presents the results of this decomposition of the MSE in the simulation exercise.

While the bias squared component of all estimates of expected returns is insignificant, the variance of the predictive regression estimates of expected returns is more than five times larger than the variance of the SOP estimates. This variance, due to estimation error, is therefore the main weakness of predictive regressions. The historical mean estimates of expected returns present a variance slightly lower than the SOP estimates. Regarding the covariance term, the SOP and predictive regressions estimates of expected return have significant positive correlations with the true expected return, which contribute to reducing the MSE. In contrast, the historical mean is negatively correlated with the true expected return, which adds to its MSE. Overall, the superior performance of the SOP method relative to the predictive regression comes from its lower variance and its higher correlation with the true expected return. The superior performance of the SOP method relative to the historical

Table 9
Monte Carlo simulation: mean square error of return forecasts.
This table presents the results of a Monte Carlo simulation of the economy in Binsbergen and Koijen (2010). The simulation generates ten thousand samples of 80 years of returns, dividend growth, and the dividend–price ratio for this economy. In each simulation of the economy, annual forecast of returns are estimated, alternatively, under the historical mean, predictive regression with the log dividend–price ratio as conditioning variable, and sum-of-the-parts (SOP) with no multiple growth methods using only past data. The forecast errors are given by the difference between the return forecasts and the true expected returns from the simulation. Panel A reports the mean, median, and other percentiles (across simulations) of the root mean square errors (RMSE) of each method. Panel B reports each component of the mean square errors (MSE) decomposition. Bias square of estimator is given by $(1/(T-s_0))\sum_{s=s_0}^{T-1}[\mathrm{E}(\hat{\mu}_s-\mu_s)]^2$. Variance of estimator is given by $(1/(T-s_0))\sum_{s=s_0}^{T-1}\mathrm{Var}(\hat{\mu}_s)$. Variance of true expected returns is given by $(1/(T-s_0))\sum_{s=s_0}^{T-1}\mathrm{Var}(\mu_s)$. Covariance of estimator and true expected returns is given by $(1/(T-s_0))\sum_{s=s_0}^{T-1}\mathrm{Cov}(\hat{\mu}_s,\mu_s)$.

	Historical mean	Predictive regression	SOP with no multiple growth
Panel A: Distribution of the root mean square error (× 100)			
Mean	4.94	3.73	2.87
10th percentile	3.22	2.04	1.90
25th percentile	3.85	2.62	2.23
Median	4.72	3.48	2.71
75th percentile	5.77	4.58	3.32
90th percentile	6.95	5.78	4.05
Panel B: Mean square error decomposition (× 1000)			
Square bias of estimator	0.244	0.017	0.007
Variance of estimator	0.264	2.697	0.511
Variance of true expected returns	1.951	1.951	1.951
−2 × covariance of estimator and true expected returns	0.200	−3.055	−1.567
Mean square error	2.658	1.608	0.902

mean is explained by its much higher correlation with the true expected return.

Finally, we compute out-of-sample R-squares in the simulations. For predictive regressions and the SOP method, the R-squares are 4.03% and 7.17%, respectively. These values compare with an R-square of 10.74% for the true expected return, which constitutes an upper bound for this statistic. The SOP method is close to being as efficient as the true expected return (which is unfeasible outside of simulation experiments) in predicting returns.

5. Conclusion

We propose forecasting separately the dividend–price ratio, the earnings growth, and the price–earnings growth components of stock market returns: the sum-of-the-parts method. Our method exploits the different time series properties of the components. We apply the SOP method to forecast stock market returns out of sample over 1927–2007. The SOP method produces statistically and economically significant gains for investors and performs better out of sample than the historical mean or predictive regressions. The gain in performance of the SOP method relative to predictive regressions is mainly due to the absence of estimation error.

Our results have important consequences for corporate finance and investments. The SOP forecasts of the equity premium can be used for cost of capital calculations in project and firm valuation. The results presented suggest that discount rates and corporate decisions should depend more on market conditions. In the investment world, we show that there are important gains from timing the market. To the extent that what we are capturing is excessive predictability rather than a time-varying risk premium, the success of our analysis eventually destroys its usefulness. Once enough investors follow our approach to predict returns, they will impact market prices and again make returns unpredictable.

Appendix A. Shrinkage approach

Following Connor (1997), we transform the estimated coefficients of the predictive regression in Eq. (2) by

$$\beta^* = \frac{s}{s+i}\hat{\beta} \tag{32}$$

and

$$\alpha^* = \overline{r}_s - \beta^* \overline{x}_s,$$

where $\overline{x}_s$ is the historical mean of the predictor up to time s. In this way, the slope coefficient is shrunk toward zero, and the intercept changes to preserve the unconditional mean return. The shrinkage intensity i can be thought of as the weight given to the prior of no predictability. It is measured in units of time periods. Thus, if i is set equal to the number of data periods in the data set s, the slope coefficient is reduced by half. Connor (1997) shows that it is optimal to choose $i = 1/\rho$, where ρ is the expectation of a function of the regression R-square:

$$\rho = \mathrm{E}\left(\frac{R^2}{1-R^2}\right) \approx \mathrm{E}(R^2). \tag{33}$$

This is the expected explanatory power of the model. We use i=100 with yearly data and i=1,200 with monthly data. This would give a weight of one hundred years of data to the prior of no predictability. Alternatively, we can interpret this as an expected R-square of approximately 1% for predictive regressions with yearly data and less than 0.1% with monthly data, which seems reasonable in light of findings in the literature. This means that if we run the predictive regression with 30 years of data, the slope coefficient is shrunk to 23% [= 30/(100+30)] of its estimated size.

Finally, we use these coefficients to forecast the stock market return r as

$$\hat{\mu}_s = \alpha^* + \beta^* x_s. \tag{34}$$

As in the predictive regression approach in Eq. (32)– (33), we apply shrinkage to the estimated coefficients of the multiple growth regression in Eq. (17):

$$\beta^* = \frac{s}{s+i}\hat{\beta} \tag{35}$$

and

$$\alpha^* = -\beta^* \overline{x}_s, \tag{36}$$

which makes the unconditional mean of the multiple growth equal to zero.

We also apply shrinkage to the estimated coefficients of the regression of the realized multiple growth on the expected multiple growth in Eq. (21) as follows:

$$d^* = \frac{s}{s+i}\hat{d} \tag{37}$$

and

$$c^* = -d^*(-\overline{\hat{u}}_s), \tag{38}$$

$$c^* = d^* \overline{\hat{u}}_s, \tag{39}$$

where $\overline{\hat{u}}_s$ is the sample mean of the regression residuals up to time s (not necessarily equal to zero). This assumes that the unconditional expectation of the multiple growth is equal to zero. That is, with no information about the state of the economy, we do not expect the multiple to change.

Appendix B. Definition of predictors of stock returns

Stock variance (SVAR): Sum of squared daily stock market returns on the S&P 500.

Default return spread (DFR): Difference between long-term corporate bond and long-term bond returns.

Long-term yield (LTY): Long-term government bond yield.

Long-term return (LTR): Long-term government bond return.

Inflation (INFL): Growth in the consumer price index with a one-month lag.

Term spread (TMS): Difference between the long-term government bond yield and the T-bill.

Treasury bill rate (TBL): Three-month Treasury bill rate.

Default yield spread (DFY): Difference between BAA- and AAA-rated corporate bond yields.

Net equity expansion (NTIS): Ratio of 12-month moving sums of net issues by NYSE-listed stocks to NYSE market capitalization.

Return on equity (ROE): Ratio of 12-month moving sums of earnings to book value of equity for the S&P 500.

Dividend payout ratio (DE): Difference between the log of dividends (12-month moving sums of dividends paid on S&P 500) and the log of earnings (12-month moving sums of earnings on S&P 500).

Earnings price ratio (EP): Difference between the log of earnings (12-month moving sums of earnings on S&P 500) and the log of prices (S&P 500 index price).

Smooth earnings price ratio (SEP): Ten-year moving average of earnings–price ratio.

Dividend price ratio (DP): Difference between the log of dividends (12-month moving sums of dividends paid on S&P 500) and the log of prices (S&P 500 index price).

Dividend yield (DY): Difference between the log of dividends (12-month moving sums of dividends paid on S&P 500) and the log of lagged prices (S&P 500 index price).

Book-to-market (BM): Ratio of book value to market value for the Dow Jones Industrial Average.

References

Ang, A., Bekaert, G., 2007. Stock return predictability: Is it there? Review of Financial Studies 20, 651–707.

Arnott, R., Bernstein, P., 2002. What risk premium is normal? Financial Analysts Journal 58, 64–85.

Ashley, R., 2006. Beyond optimal forecasting. Unpublished working paper, Virginia Polytechnic Institute and State University.

Baker, M., Wurgler, J., 2000. The equity share in new issues and aggregate stock returns. Journal of Finance 55, 2219–2257.

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

Bogle, J., 1991a. Investing in the 1990s. Journal of Portfolio Management 17, 5–14.

Bogle, J., 1991b. Investing in the 1990s: Occam's razor revisited. Journal of Portfolio Management 18, 88–91.

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

Boudoukh, J., Michaely, R., Richardson, M., Roberts, M., 2007. On the importance of measuring payout yield: implications for empirical asset pricing. Journal of Finance 62, 877–915.

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

Brandt, M., 2009. Portfolio choice problems. In: Ait-Sahalia, Y., Hansen, L.P. (Eds.), Handbook of Financial Econometrics. Elsevier Science, Amsterdam.

Brandt, M., Kang, Q., 2004. On the relationship between the conditional mean and volatility of stock returns: a latent VAR approach. Journal of Financial Economics 72, 217–257.

Breen, W., Glosten, L., Jagannathan, R., 1989. Economic significance of predictable variations in stock index returns. Journal of Finance 64, 1177–1189.

Burnside, C., Eichenbaum, M., Kleshchelski, I., Rebelo, S. Do peso problems explain the returns to the carry trade? Review of Financial Studies, forthcoming.

Burnside, C., Eichenbaum, M., Rebelo, S., 2007. The returns to currency speculation in emerging markets. American Economic Review Papers and Proceedings 97, 333–338.

Campbell, J., 1987. Stock returns and term structure. Journal of Financial Economics 18, 373–399.

Campbell, J., 2008. Estimating the equity premium. Canadian Economic Review 41, 1–21.

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

Campbell, J., Thompson, S., 2008. Predicting the equity premium out of sample: Can anything beat the historical average? Review of Financial Studies 21, 1509–1531.

Campbell, J., Vuolteenaho, T., 2004. Inflation illusion and stock prices. American Economic Review 94, 19–23.

Campbell, J., Yogo, M., 2006. Efficient tests of stock return predictability. Journal of Financial Economics 81, 27–60.

Cavanagh, C., Elliott, G., Stock, J., 1995. Inference in models with nearly integrated regressors. Econometric Theory 11, 1131–1147.

Clark, T., McCracken, M., 2001. Test of equal forecast accuracy and encompassing for nested models. Journal of Econometrics 105, 85–110.

Claus, J., Thomas, J., 2001. Equity premia as low as 3%? Evidence from analysts' earnings forecasts for domestic and international stock markets. Journal of Finance 56, 1629–1666.

Cochrane, J., 2008. The dog that did not bark: a defense of return predictability. Review of Financial Studies 21, 1533–1575.

Connor, G., 1997. Sensible return forecasting for portfolio management. Financial Analysts Journal 53, 44–51.

Diebold, F., Mariano, R., 1995. Comparing predictive accuracy. Journal of Business and Economic Statistics 13, 253–263.

Dow, C., 1920. Scientific stock speculation. Magazine of Wall Street.

Fama, E., French, K., 1988. Dividend yields and expected stock returns. Journal of Financial Economics 22, 3–25.

Fama, E., French, K., 1998. Value versus growth: the international evidence. Journal of Finance 53, 1975–1999.

Fama, E., French, K., 2002. The equity premium. Journal of Finance 57, 637–659.

Fama, E., Schwert, G.W., 1977. Asset returns and inflation. Journal of Financial Economics 5, 115–146.

Ferson, W., Sarkissian, S., Simin, T., 2003. Spurious regressions in financial economics. Journal of Finance 58, 1393–1413.

Foster, F.D., Smith, T., Whaley, R., 1997. Assessing goodness-of-fit of asset pricing models: the distribution of the maximal R^2. Journal of Finance 52, 591–607.

French, K., Schwert, G.W., Stambaugh, R., 1987. Expected stock returns and volatility. Journal of Financial Economics 19, 3–30.

Ghysels, E., Santa-Clara, P., Valkanov, R., 2005. There is a risk-return trade-off after all. Journal of Financial Economics 76, 509–548.

Goyal, A., Santa-Clara, P., 2003. Idiosyncratic risk matters!. Journal of Finance 58, 975–1007.

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

Guo, H., 2006. On the out-of-sample predictability of stock market returns. Journal of Business 79, 645–670.

Hodrick, R., 1992. Dividend yields and expected stock returns: alternative procedures for inference and measurement. Review of Financial Studies 5, 357–386.

Ibbotson, R., Chen, P., 2003. Long-run stock returns: participating in the real economy. Financial Analysts Journal 59, 88–98.

Inoue, A., Kilian, L., 2004. In-sample or out-of-sample tests of predictability: Which one should we use? Econometric Reviews 23, 371–402.

Kothari, S.P., Shanken, J., 1997. Book-to-market, dividend yield, and expected market returns: a time series analysis. Journal of Financial Economics 44, 169–203.

Lamont, O., 1998. Earnings and expected returns. Journal of Finance 53, 1563–1587.

Lettau, M., Ludvigson, S., 2001. Consumption, aggregate wealth, and expected stock returns. Journal of Finance 56, 815–849.

Lewellen, J., 2004. Predicting returns with financial ratios. Journal of Financial Economics 74, 209–235.

McCracken, M., 2007. Asymptotics for out of sample tests of Granger causality. Journal of Econometrics 140, 719–752.

Meese, R., Rogoff, K., 1983. Empirical exchange rate models of the seventies: Do they fit out of sample? Journal of International Economics 14, 3–24.

Nelson, C., 1976. Inflation and rates of return on common stocks. Journal of Finance 31, 471–483.

Nelson, C., Kim, M., 1993. Predictable stock returns: the role of small sample bias. Journal of Finance 48, 641–661.

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

Pontiff, J., Schall, L., 1998. Book-to-market ratios as predictors of market returns. Journal of Financial Economics 49, 141–160.

Ritter, J., Warr, R., 2002. The decline of inflation and the bull market of 1982–1999. Journal of Financial and Quantitative Analysis 37, 29–61.

Stambaugh, R., 1999. Predictive regressions. Journal of Financial Economics 54, 375–421.

Valkanov, R., 2003. Long-horizon regressions: theoretical results and applications. Journal of Financial Economics 68, 201–232.

Welch, I., 2000. Views of financial economists on the equity premium and other issues. Journal of Business 73, 501–537.