Generalized financial ratios to predict the equity premium[☆]

Andres Algaba[a,*], Kris Boudt[a,b]

[a] *Solvay Business School, Vrije Universiteit Brussel, Belgium*
[b] *Faculty of Economics and Business, Vrije Universiteit Amsterdam, The Netherlands*

Abstract

Empirical evidence for the price-dividend ratio to be a predictor of the equity premium is weak. We argue that changes in the economic conditions and market composition lead to a time-varying relationship between prices, dividends and the equity premium. Exploiting the information in the rolling window log-log regression of stock prices on dividends, we obtain the Generalized Price-Dividend Ratio (GPDR), that compares the price per share with a time-varying transformation of the dividend per share. The GPDR leads to economic and statistical gains when forecasting the equity premium of the S&P 500 at the 1, 3, 6 and 12 month horizon, as compared to using the classical price-dividend ratio or the prevailing historical average excess market return. Similar improvements are obtained for Generalized Financial Ratios based on the corporate earnings and book value.

Keywords: Equity premium, ERP, Forecast combination, Price-dividend ratio, Financial ratios, Time-varying parameters
JEL: C10, G11

[☆] We thank Sushanta Mallick (the Editor) and three anonymous referees, as well as Sophie Béreau, Christopher J. Neely, Bilel Sanhaji, Stefan Straetmans, Steven Vanduffel, Marjan Wauters, participants at the VVE Day (Antwerp 2016), the ITISE conference (Granada 2016), the Solvay ES Research Day (Brussels 2016), the 6th PhD Student Conference in International Macroeconomics and Financial Econometrics (Paris 2017) and the SoFIE Summer School (Brussels 2017) for their valuable comments. This research did not receive any specific grant from funding agencies in the public, commercial, or not-for-profit sectors. The views expressed in this paper are the sole responsibility of the authors. Any remaining errors or shortcomings are those of the authors.
[*] Corresponding author. Vrije Universiteit Brussel, Pleinlaan 2, 1050 Brussel, Belgium.
Email addresses: andres.algaba@vub.be (Andres Algaba), kris.boudt@vub.be (Kris Boudt)

"Clearly, Graham had found a good indicator of relative value over the near term. But he had also found yet one more indication of how valuation standards can change with the times."

- John C. Bogle, 2005.

1. Introduction

Risk and reward go hand in hand when deciding to invest in equities. When used as an ex-ante concept, the equity premium expresses the reward (or premium) an investor requires for taking on additional risk by investing in a risky equity portfolio compared to investing in a risk-free asset, such as an AAA-rated government bond. It is also called the equity risk premium to emphasize that it is a compensation for the risk taken when investing in stocks. Graphically, it corresponds to the slope of the capital allocation line connecting the risk-free asset with the risky equity portfolio in the mean-standard deviation graph of investment opportunities. Ex-post, the realized equity premium is computed as the observed excess return of a stock portfolio over the risk-free rate. Welch (2000) refers to the equity premium as "perhaps the single most important number in financial economics". A reliable prediction of the equity premium is particularly important for asset allocation and the valuation of risky assets.

In this paper, we contribute to the literature that uses financial ratios to predict the equity premium over a horizon that is in between one and twelve months[1,2]. This horizon is relevant for the tactical component in market timing the asset allocation between the risky (equity) portfolio and the risk-free asset. It features prominently in both academic research (see e.g., Brennan et al., 1997; Faber, 2007; Neely et al., 2014) and in real-life asset allocation solutions such as the Hull Tactical US ETF (Hull and Qiao, 2017).

[1] Other strands of the literature study the forecasting of the equity premium at both shorter and longer horizons. We refer to the work of Chan et al. (1992) for an analysis of the equity premium at the daily horizon, to Duarte and Rosa (2015) for the monthly, quarterly and annual horizon (which overlap with the ones we consider), to Graham and Harvey (2005) and Welch (2000) for an overview of most of the relevant longer horizons, to Prat (2013) for a comparison between the short- and long-term equity premium, and to Damodaran (2009) for the determinants of the equity premium, different approaches to measure it, a comparison between the equity premium in different countries and the relationship between the equity premium and other risk premiums (i.e. in the bond market). The longer horizons are relevant for, among others, the valuation of equity investments and strategic asset allocation, while the short horizons can be used for market timing the investment decision.

[2] Another important part of the equity premium literature is the research on the existence of the so-called equity premium puzzle. Stocks should generate higher returns on average than bonds, because they are riskier. The puzzle is that the excess returns of stocks over bonds are too high to be explained by risk alone. For a recent overview of the equity premium puzzle, we refer the interested reader to Cochrane (2017) and the references therein.

It is important to stress that the point-in-time value of a financial ratio is not designed to forecast the equity premium. Instead, it is a valuation tool to express the market value in relative terms compared to (functions of) accounting-based measures such as the dividend or earnings per share. We find that transformations of the financial ratios that properly account for the time-varying relationship between the variables at stake are better suited for predicting the equity premium at the intra-year horizon than the original financial ratios.

Within the large family of financial ratios, we focus on the price-dividend ratio since dividends play such a direct role in the formation of the equity premium. We then extend the approach to other well-known financial ratios, such as the (cyclically adjusted) price-earnings ratio, the price-earnings to growth ratio, the price-to-book ratio and the bond-equity yield ratios.

Historically, the aggregate price-dividend ratio, comparing the value weighted average share price with the corresponding dividend value, has been pivotal to obtain predictions of the equity premium. The intuition behind this is that, all other things being equal, a historically high (low) value of the price-dividend ratio implies that the stock (market) is likely to be (under-) overvalued and that a correction in its values is likely to follow (Cochrane, 1997).

Despite the intuitive appeal of the channel of mean-reversion through which the ratio between prices and dividends predicts the equity premium, there is only weak empirical evidence that the price-dividend ratio is effectively a good predictor of the US equity premium[3]. In early research, Fama and French (1988), Rozeff (1984) and Shiller (1984) find that the price-dividend ratio is significantly negatively correlated with the future US equity premium, on horizons varying from one month to four years. More recently, Goyal and Welch (2003) find that all the evidence for predictive power has disappeared since 1990, and that even before, the out-of-sample performance of the price-dividend ratio is rather poor. In particular, Goyal and Welch (2008) find that, for US equities over the period 1871–2005, the price-dividend ratio is unable to outperform the prevailing historical average equity premium as a predictor of the equity premium. However, Cochrane (2008) claims that the time-variation of the price-dividend ratio must contain some (excess) return predictability

[3]The US equity premium is defined as the total return on the US stock market minus the corresponding prevailing risk-free rate. The value-weighted Standard and Poor's Composite Stock Price Index (S&P 500) is most often used as a proxy for the US equity universe. As the risk-free rate, we use the total return of investing in the 3-month US Treasury Bill over the same horizon as the equity investment

as he finds powerful statistical evidence that it does not predict dividend growth.

We conjecture that the weak predictive power of the price-dividend ratio for the equity premium comes from its failure to capture two important effects on the equity premium. The first failure is that the price-dividend ratio does not account for the time-varying economic conditions. This issue can be related to traditional pricing theory, such as Gordon's (1959) dividend discount model. This model discounts the expected future dividends to calculate a present value for the stock. The discount factor includes the constant expected required rate of return (r) and the constant expected dividend growth (g). The price-dividend ratio fails to incorporate this discounting information. We illustrate this in Figure 1 where we plot the monthly S&P 500 price-dividend ratio, the two-year simple moving average of the S&P 500 pay-out ratio and the long-term interest rate from January 1956 until December 2014. The long-term interest rate is the 10-year US Treasury constant maturity rate. As expected, we find a negative association between interest rates and the price-dividend ratio. This provides only a partial explanation of the the price-dividend ratio. Note, e.g., that the price-dividend ratio is relatively stable in the sixties and the beginning of the seventies despite the increase in interest rates from 3.8 to 7.8%. Second, the price-dividend ratio at the end of sample is almost double the one at the beginning of the sample, notwithstanding that they have in common that for both periods the interest rate is historically low.

The second failure is that the price-dividend ratio does not account for the time-varying market structure. It has the implicit assumption that the composition of the prices and dividends and their relationship never changes. This is not true, specifically not for a value-weighted index like the S&P 500, since its composition changes daily. Using the same price-dividend ratio would imply comparing the market portfolio of 1980, consisting mostly of energy companies such as Shell Oil, Standard Oil of California and Standard Oil of Indiana, with the market portfolio of 2015, consisting mostly of IT companies such as Apple, Facebook and Google. The market changes and so should the dynamic properties of the ratio. This is connected to the observation in Fama and French (2001) that, due to the new listings of low profit/high growth opportunity firms, the average dividend paid by S&P 500 firms compared to the market value has reduced dramatically in the past decades. Moreover, a lot of companies pursue different dividend policies to maximize their market value (Al-Malkawi et al., 2014). Figure 1 shows the declining long-term interest rate since the eighties which is one of the possible reasons for the recent changes in the composition of equity in-

Figure 1: Time series of the monthly S&P 500 price-dividend ratio, the two-year simple moving average of the S&P 500 pay-out ratio and long-term interest rate over the period January 1956 - December 2014.

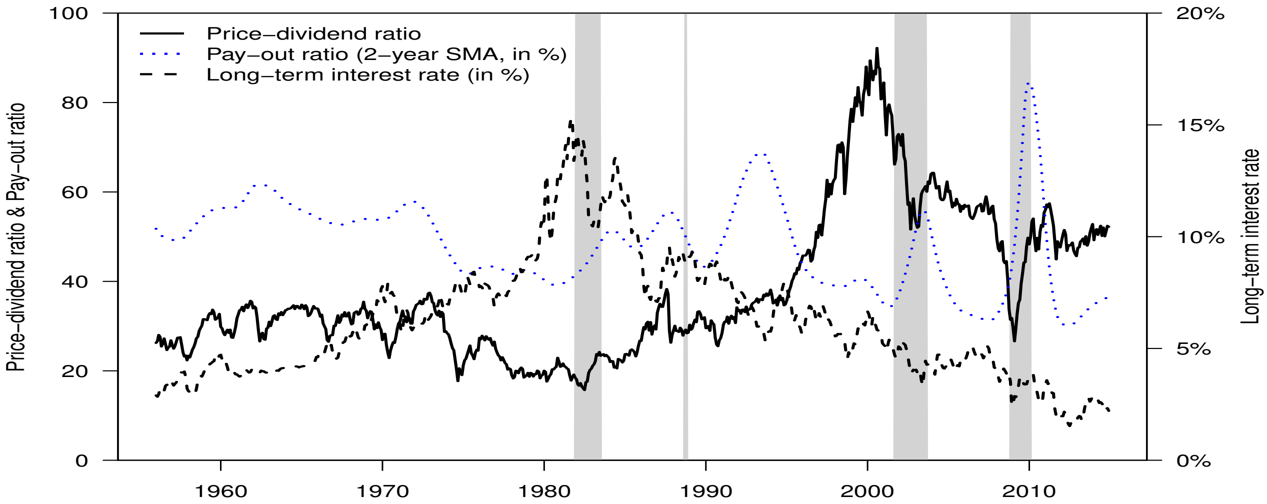

Note: This figure contains monthly observations of the S&P 500 price-dividend ratio (prices divided by dividends), 2-year simple moving average of the S&P 500 pay-out ratio (dividends divided by earnings) and the long-term interest rate. Dividends (earnings) are computed as the rolling sum of the dividends (earnings) of the last 12 months. The long-term interest rate is the 10-year US Treasury constant maturity rate (GS10). For a more detailed description of the data, we refer to Section 3 and 6. Shaded regions indicate the four largest drawdown periods of the S&P 500 during the out-of-sample period (01/1976 – 12/2014), namely 12/1980 – 07/1982 (the early 1980s recession in the US), 09/1987 – 11/1987 (Black Monday), 09/2000 – 09/2002 (the collapse of the dot-com bubble), 11/2007 – 02/2009 (the financial crisis).

dexes, such as the S&P 500. It further shows the declining trend and the wide variability in the S&P 500 pay-out ratio. These changes in the market's dividend policy may affect the elasticity from prices to dividends. This elasticity can also be affected on the investor's side via changes (structural or not) in investors' attitudes towards dividends and taxes (Polimenis and Neokosmidis, 2016).

The methodological innovation of this paper consists of solving both failures by comparing the equity's market value with a non-linear transformation of the dividend per share, forming the Generalized Price-Dividend Ratio (GPDR). Importantly, the parameters in the GPDR's functional form are time-varying. We propose to capture this time-variation by the use of Nadaraya-Watson (Nadaraya, 1964; Watson, 1964) weighted parameter estimates of the log-log regression of the stock prices on its corresponding dividends over rolling estimation samples. We consider various weighting methods and relate these to the business cycle. Besides the univariate GPDR forecasting equation, we use two types of forecast combinations. The first one combines GPDRs computed over various estimation windows. It uses least squares optimized weights under the form of a Heterogeneous AutoRegressive, or HAR, forecasting equation to account for the heterogeneity in the market reaction to price changes (Corsi, 2009). The second approach combines the GPDR-based equity premium forecast with the prevailing historical average equity premium estimate by means of a switching forecasting equation. The switches are driven by past relative performance. All the proposed methods are in explicit form and thus are simple and fast to compute.

In an extensive out-of-sample evaluation for the period of January 1976 until December 2014, we find that the use of the GPDR leads to economic and statistical gains in predicting the monthly S&P 500 index equity premium compared to the classical price-dividend ratio, and the prevailing historical average, which has been shown to be a strong benchmark (Goyal and Welch, 2008). We further use Goyal and Welch's (2003, 2008) recursive residuals (out-of-sample) graphical approach to show that the good performance is not bounded to any specific subperiod. We also find that the relative forecasting outperformance holds at the quarterly, semi-annual and annual horizon.

We use the same economic intuition that motivates the use of the GPDR to generalize other financial ratios. Instead of dividends, we use other proxies for the fundamental characteristic of a stock, such as the corporate earnings and book value per share. The framework of generalizing financial ratios by including time-varying parameters is thus flexible and leads to the family of the Generalized Financial Ratios.

The bottom line of our research is that, for forecasting the equity premium at the monthly horizon, the generalized financial ratio leads to more accurate predictions than when the traditional financial ratio is used. This result is useful for both academics and practitioners who need an equity premium forecast when optimizing risk and reward of their investment portfolio.

The remainder of the paper is organised as follows. In Section 2 the Generalized Price Dividend Ratio (GPDR) is introduced. The data is described in Section 3. Sections 4 and 5 present the out-of-sample evaluation of the gains in using the GPDR instead of the classical price-dividend ratio when forecasting the equity premium. Section 6 extends the GPDR to a class of Generalized Financial Ratios, including also the generalized (cyclically adjusted) price-earnings ratio, price-earnings to growth ratio, price-to-book ratio and bond-equity yield ratios. Section 7 concludes.

2. Methodology

In this section, we first set our notation. We then introduce the main components of the GPDR framework to predict the equity premium: (i) general definition of the GPDR; (ii) Nadaraya-Watson estimates of the time-varying parameters in the GPDR; (iii) the single-regime univariate and HAR-GPDR, and switching forecasting equations to predict the equity premium. We further show that the approach of forecasting the equity premium with the GPDR nests the special cases of forecasting the equity premium with the classical and modified price-dividend ratio. Finally, we explain that the GPDR is at the intersection of fundamental and technical analysis.

2.1. Notation

For convenience in notation, we present the framework in terms of forecasting the monthly equity premium. We then apply it to forecasting the equity premium not only at the one-month horizon, but also the quarterly, semi-annual and annual frequency using trailing 3, 6 and 12-month returns.

Throughout the paper, we use P_t and D_t to denote the end-of-period t stock prices and dividends, whereby the observation frequency considered is monthly. The proposed framework is generally applicable to forecasting (excess) returns on assets at different levels of granularity: individual equities, sector indices and market-wide equity indices. As in Goyal and Welch (2008), we apply the proposed methodology to the US equity market, for which the excess return on the market capitalization weighted S&P 500 portfolio is called the US equity premium (*ERP*). It is computed as follows:

$$ERP_t = \log(1 + r_t^m) - \log(1 + r_t^f), \tag{1}$$

where r_t^f denotes the (prevailing) US Treasury Bill rate and r_t^m is the total return (capital and dividend gains combined) on the S&P 500, *i.e.* $r_t^m = [P_t - P_{t-1} + D_t]/P_{t-1}$, where P_t is the S&P 500 price index value and D_t are the dividends gained during the return period.

2.2. Definition of the GPDR

The framework of forecasting the equity premium using the GPDR is a generalization of the approach that predicts the equity premium using the classical price-dividend ratio. The theoretical foundation for mean-reversion and thus the predictive power of the classical price-dividend ratio comes from the cointegration relationship between log-prices (p) and log-dividends (d). Campbell and Shiller's (1988) log-linear dividend-ratio model predicts this relationship. Adding a time index we obtain the following linear equation:

$$p_t = \alpha + \beta d_t + \varepsilon_t, \tag{2}$$

where ε_t is assumed to be a stationary error term with zero mean, α is a constant and β is the cointegration coefficient between log-prices and log-dividends. If we exponentiate both sides of the equation, we get the following fundamental price equation under Campbell and Shiller's (1988) log-linear approximation:

$$\begin{aligned} P_t &= \lambda D_t^{\beta} \exp(\varepsilon_t), \\ \text{with } \lambda &= \exp(\alpha). \end{aligned} \tag{3}$$

Since we assume the expected value for ε_t to be equal to zero, the expected value for the market price is the following:

$$\mathbb{E}[P_t | \varepsilon_t = 0] = \lambda D_t^{\beta}. \tag{4}$$

The proposed "Generalized Price-Dividend Ratio" (GPDR) is then immediately obtained by dividing the price by λD_t^{β}. The GPDR is thus defined as follows:

$$GPDR_t(\lambda_t, \beta_t) = \frac{P_t}{\lambda_t D_t^{\beta_t}}, \tag{5}$$

where λ_t and β_t are parameters that are allowed to be time-varying. Note that the GPDR is essentially the observed market price divided by its theoretical value under the model in Equation (4). It thus follows that a GPDR value of around one indicates that the stock market is correctly priced, whereas a GPDR lower than unity means that the stock market is underpriced and vice versa.

The time-varying parameters aim to overcome two shortcomings that the classical price-dividend ratio has. The first shortcoming concerns the information about the time-varying discount factor, the risk appetite, the expectations about the interest rate and the dividend growth, among others. All this information about the time-varying economic conditions is currently ignored when using the classical-price dividend ratio as a valuation indicator. We do not explicitly model the discount factor that we use in our GPDR, instead we estimate it via a data-driven approach. Our approach can be reconciled with most of the commonly used theoretical discount factor models. As an example, we use Gordon's (1959) dividend discount model to show how any discount model is compatible with the GPDR[4]. Gordon's (1959) dividend discount model looks as follows:

$$P_t = kD_t, \quad \text{with } k = \frac{1+g}{r-g}. \tag{6}$$

The model is actually a present value model that discounts all the future cashflows (in the form of dividends) in order to give a present value of the stock. The discount factor (k) in this model contains the constant expected required rate of return (r) and the constant expected dividend growth (g). If we compare Equation (4) with Equation (6), one can see that the denominator of the GPDR can be considered as a generalized present value model. The parameter λ can then be interpreted as the estimated discount factor. The classical price-dividend ratio does not include any information concerning the discount factor, as it is equal to unity. In our GPDR we add the information that is concealed in the discount factor by allowing for a time-varying discount factor λ_t.

The second shortcoming is about the market's changing composition, varying dividend policy and the shifting of investors' attitudes towards dividends and taxes. All these elements may affect the elasticity from prices to dividends. The classical price-dividend ratio does, by design, not account for this elasticity and its changes. The cointegration coefficient (β) can be interpreted as an elasticity coefficient, due to the log-log regression in Equation (2). Comparing Equation (4) and Equation (6), the cointegration coefficient is assumed to be equal to unity in the classical price-dividend ratio. We add the elasticity information in our GPDR by using the time-varying cointegration coefficient β_t.

2.3. Estimation of the time-varying parameters

We have now defined the general functional form of the GPDR, for which the value depends on the α_t and β_t parameters corresponding to Equation (2). It is important to remember that the economic intuition for the time-variation in the parameters is that it is partly driven by the time-variation in other variables that are not explicit in the definition of the GPDR, like the risk appetite, the level of risk aversion and the value of the interest rate. Since our focus is on forecasting the ERP, understanding the exact nature of the dependence of α_t and β_t on those confounding factors is of secondary importance. Moreover, the functional

[4] Examples of alternative models are those with a time-varying required rate of return (Campbell and Shiller, 1988) or a present value model based on consumption (LeRoy, 1973; Lucas, 1978).

relationship is likely to be time-varying. We therefore take an agnostic view on the drivers of the α_t and β_t parameters and propose to estimate these parameters using semi-parametric estimation techniques.

The proposed semi-parametric updating scheme uses a weighted least squares approach to estimate α_t and β_t in the cointegration equation regressing log-prices on log-dividends in Equation (2). The weight functions are asymmetric to avoid look-ahead bias and use decaying weights to reflect that the most recently observed values are more informative about the current fundamentally weighted value of the GPDR than the more remote observations. The resulting estimates belong to the class of Nadaraya-Watson estimators (Nadaraya, 1964; Watson, 1964). They are given by:

$$\hat{\beta}_t = \frac{\sum_{s=1}^{t} w(t-s)(p_s - \hat{p}_t)(d_s - \hat{d}_t)}{\sum_{s=1}^{t} w(t-s)(d_s - \hat{d}_t)^2}, \quad \hat{\alpha}_t = \hat{p}_t - \hat{\beta}_t \hat{d}_t, \tag{7}$$

where:

$$\hat{p}_t = \sum_{s=1}^{t} w(t-s) p_s, \quad \hat{d}_t = \sum_{s=1}^{t} w(t-s) d_s. \tag{8}$$

The function w is a generic notation for the weight function, that needs to satisfy the constraint that, for all u, $w(u) \geq 0$ and increasing, and that, for $u < 0$, $w(u) = 0$.

In this application, we will consider the weight functions that correspond to three popular methods in practice, namely estimation from inception, simple moving average estimation and the exponentially weighted moving average approach. The corresponding weight functions, where we use l to denote the time lag between the current date t and the time of observation s, are given by:

$$\begin{aligned} w^{INC}(l) &= I[0 \leq l]/t, \\ w^{SMA}(l;m) &= I[0 \leq l < m]/m, \\ w^{EWMA}(l;m;\tau) &= \tau^l I[0 \leq l < m] / \sum_{l=0}^{m-1} \tau^l. \end{aligned} \tag{9}$$

The parameter m determines the length of the rolling estimation window, $I[\cdot]$ is an indicator function and $\tau > 0$ is the decay coefficient. Note that when w^{INC} is used, the obtained estimates are the standard ordinary least squares (OLS) coefficients on the sample available from inception. Every observation is equally weighted and all the observations, from the starting point until the last observation (at time t), are used. The second weighting method (w^{SMA}) is the rolling window estimation method, which can be considered to be a simple moving average. The length of the estimation window is determined via m. The third weighting method (w^{EWMA}) uses an exponentially weighted moving average to do the time aggregation on a rolling sample of m observations. The decay in the weights is determined by the strictly positive coefficient τ ensuring that recent observations carry a higher weight. The corresponding GPDRs for each weighting scheme are denoted as INC-GPDR, SMA-GPDR and EWMA-GPDR.

In order to use the SMA-GPDR and EWMA-GPDR, the values for the parameters m and τ need to be set. An economically meaningful calibration is obtained by matching the coefficients α_t and β_t with what they represent. The coefficient α_t is used to capture the time-varying discount factor λ_t, which should move with changing expectations about r and g throughout the business cycle. The elasticity of prices to dividends also depends on risk aversion and thus varies over the business cycle length. Based on this, we recommend to set the window length (m) so that it approximates the length of an average business cycle. Jadevicius and Hutson (2014) distinguish four different types of business cycles, namely the Kitchin cycle (5 years), the Julgaler cycle (9 years), the Kuznets cycle (20 years) and the Kondratieff wave (50 years). Since we allow α_t and β_t to be possibly rapidly time-varying parameters, the Kitchin cycle and the Julgaler cycle seem like natural choices. This leads to GPDR implementations using three different estimation windows, namely the expanding window from inception and the rolling window of length equal to the Kitchin cycle, and the Julgaler cycle. We further set the decay coefficient (τ) such that the half-life is equal to half the length of the window.

2.4. Forecasting equations

We consider the traditional univariate forecasting equation to use the GPDR for predicting the Equity Risk Premium (*ERP*), as well as two types of forecast combinations. We refer to these forecast combinations as the HAR-GPDR and the switching GPDR-based predictions of the *ERP*. The former uses least squares optimized weights to combine the predictive power of GPDRs computed over windows of different lengths. The latter combines the GPDR-based *ERP* forecasts with the forecasts based on the prevailing historical average *ERP* using relative past forecast performance driven weights[5].

[5] We refer to Stock and Watson (2004) and Genre et al. (2013) for a general overview on the use of forecast combinations in economics and finance.

2.4.1. Traditional single-regime univariate forecasting equation

The first functional form that we use to forecast the *ERP* using the log-transformed GPDR is widespread in the equity premium literature, namely the regression of the *ERP* on a one-period lagged predictor and a constant (see e.g., Goyal and Welch, 2008; Gupta and Mampho, 2013). If we denote the estimated $GPDR_t$ as $\widehat{GPDR}_t$, we obtain the following forecasting equation:

$$ERP_t = \delta_t + \theta_t \log(\widehat{GPDR}_{t-1}) + e_t, \tag{10}$$

where δ_t and θ_t are the time-varying intercept and slope parameter, and e_t is the error term. Note that we expect θ_t to be a negative coefficient such that a higher value of the estimated GPDR implies, *ceteris paribus*, lower future (excess) returns[6]. We use an expanding estimation window to forecast the *ERP* for out-of-sample purposes. The corresponding predicted *ERP* is then:

$$\widehat{ERP}^{\text{GPDR}}_{t+1|t} = \hat{\delta}_t + \hat{\theta}_t \log(\widehat{GPDR}_t), \tag{11}$$

where $\hat{\delta}_t$ and $\hat{\theta}_t$ are the OLS estimates of δ_t and θ_t in Equation (10) over the available estimation window.

2.4.2. Single-regime HAR-GPDR forecasting equation

In the following, we consider an extension of the univariate forecasting equation to predict the *ERP* to a forecasting model that includes multiple estimates of the GPDR computed on windows of different lengths. In analogy to the Heterogeneous AutoRegressive Realized Volatility (HAR-RV) model of Corsi (2009), we call this the HAR-GPDR model. The economic motivation is that the *ERP* is a result of the demand of heterogeneous investors. In addition to e.g., differences in endowments and risk profiles, they differ in terms of recall of market events, because of differences in age (Korniotis and Kumar, 2011), experiences (Kaustia and Knüpfer, 2008) and personal beliefs and preferences (Kuhnen and Knutson, 2011), among others. Moreover, the distribution of recall across investors is time-varying. A practical implication of this time-varying heterogeneity is that the appropriate choice of estimation window size in Section 2.3 is likely to be time-varying.

Simplifying reality, we consider that the heterogeneity leads to two possible implemen-

[6]Our primary interest is in predicting the *ERP* for which we need point estimates of δ_t and θ_t. We are not interested in the distributional properties of the estimates of δ_t and θ_t, which would require assumptions on the stationarity of the GPDR. Note however that the economic rationale of forecasting the equity premium using the GPDR is that the log-transformed GPDR is the residual of the estimated cointegration relationship between log-prices and log-dividends.

tations of the GPDR, namely the GPDR with a medium-term memory (Kitchin cycle) and with a long-term memory (Julgaler cycle). As intended in the original HAR-RV model we only use simple moving averages, and thus the SMA-GPDRs and not the EWMA-GPDRs. This leads to the following forecasting equation:

$$ERP_t = \eta_t + \theta_t^{(\mathrm{MT})} \log(\widehat{GPDR}_{t-1}^{(\mathrm{MT})}) + \theta_t^{(\mathrm{LT})} \log(\widehat{GPDR}_{t-1}^{(\mathrm{LT})}) + e_t, \tag{12}$$

where η_t is a constant, $\theta_t^{(\mathrm{MT})}$ denotes the OLS parameter for the medium-term memory GPDR and $\theta_t^{(\mathrm{LT})}$ denotes the OLS parameter for the long-term memory GPDR. This model is linear in the regression parameters and thus lends itself to convenient estimation by OLS. The resulting *ERP* prediction under the HAR-GPDR approach is then:

$$\widehat{ERP}_{t+1|t}^{\text{HAR-GPDR}} = \hat{\eta}_t + \hat{\theta}_t^{(\mathrm{MT})} \log(\widehat{GPDR}_t^{(\mathrm{MT})}) + \hat{\theta}_t^{(\mathrm{LT})} \log(\widehat{GPDR}_t^{(\mathrm{LT})}), \tag{13}$$

where $\hat{\eta}_t$, $\hat{\theta}_t^{(\mathrm{MT})}$ and $\hat{\theta}_t^{(\mathrm{LT})}$ are the OLS estimates of η_t, $\theta_t^{(\mathrm{MT})}$ and $\theta_t^{(\mathrm{LT})}$ in Equation (12) over the available estimation window.

2.4.3. *Switching forecasting equation*

The linear forecasting equations in (10)–(13) have in common that they assume that the use of the GPDR improves the forecasting precision of the *ERP* compared to the use of the prevailing historical average. Such a supremacy of one model to forecast the equity premium contradicts the so-called Adaptive Market Hypothesis (AMH) of Lo (2005), which states that the informational efficiency, and thus the economic gains in using valuation indicators such as the GPDR to forecast the equity premium, are time-varying and depend on the prevailing market regime. The approach may perform well in certain environments and poorly in others. We cannot expect one variable to be always able to predict the *ERP* in all market regimes. That is why we recommend to combine the GPDR-based forecast with the prevailing historical average *ERP* forecast based on relative past forecasting performance. More precisely, we propose to predict the *ERP* using a threshold switching equation. Ideally it should switch from using the GPDR-based forecasts in (10)–(13) to the use of the prevailing historical average, when it can be expected that the latter will be more precise, and vice versa when the GPDR-forecast is expected to be more accurate.

But how to implement this in practice? First of all, we need a measure of forecasting precision. As in Goyal and Welch (2008), we use the difference in the squared prediction errors of the GPDR-based forecast and the historical-average based forecast. Secondly, we need to define how the expectation of relative performance will be defined. For this, we use a learning model inspired by the trend-following investment strategy of Faber (2007). It consists of predicting the GPDR to be outperforming over the next month, when over the past three months it has performed relatively better (i.e. lower squared forecasting errors) than over the past six months. Putting those ideas together, we obtain the following switching forecasting equation:

$$\widehat{ERP}^{\text{SW-GPDR}}_{t+1|t} = \begin{cases} \widehat{ERP}^{\text{GPDR}}_{t+1|t} & \text{if } RASE_t(6) \geq RASE_t(3) \\ \widehat{ERP}^{\text{HA}}_{t+1|t} & \text{otherwise,} \end{cases} \tag{14}$$

where $\widehat{ERP}^{\text{GPDR}}_{t+1|t}$ is the prediction from (11), and

$$\widehat{ERP}^{\text{HA}}_{t+1|t} = (1/t) \sum_{s=1}^{t} ERP_s. \tag{15}$$

The state variable uses the S-months Rolling Average Squared Error (RASE) between the GPDR-based forecast and the use of the prevailing historical average:

$$RASE_t(S) = \frac{1}{S} \sum_{s=t-S+1}^{t} (\widehat{ERP}^{\text{GPDR}}_s - ERP_s)^2 - (\widehat{ERP}^{\text{HA}}_s - ERP_s)^2. \tag{16}$$

We thus compare the RASE over the most recent six and three months. When the difference between both is negative (i.e. when the average squared prediction error of the past six months is smaller than the average squared prediction error of the past three months), the performance of the GPDR model has deteriorated compared to the past six months. The timing model then assumes to have reached a period where it is difficult to predict the ERP with the GPDR. It thus concludes that the model sophistication of the GPDR has no added value and the simple approach of using the prevailing historical average is therefore recommended. Note that the switching forecasting equation does not suffer from a look-ahead bias as only the information up to and including time t is used.

2.5. Relation to the (modified) price-dividend ratio

The GPDR is an improvement of the classical Price-Dividend Ratio (PDR) and the Modified Price-Dividend Ratio (MPDR) defined by Polimenis and Neokosmidis (2016). Both the PDR and MPDR are nested as special cases in the GPDR. To illustrate the added value of the generalization, let us first consider the log-transformation of the estimated GPDR:

$$\widehat{gpdr}_t = p_t - \hat{\alpha}_t - \hat{\beta}_t d_t. \tag{17}$$

Plugging the right hand side definition of $\log(\widehat{GPDR}_{t-1})$ in the *ERP* forecast equation of (10) yields:

$$ERP_t = \delta_t + \theta_t(p_{t-1} - \hat{\alpha}_{t-1} - \hat{\beta}_{t-1}d_{t-1}) + e_t. \tag{18}$$

The special cases of the PDR and the MPDR are obtained by restricting the time-variation in the time-varying parameter estimates $\hat{\alpha}_t$ and $\hat{\beta}_t$. Firstly, if we restrict both $\hat{\alpha}_t$ and $\hat{\beta}_t$ to be time-invariant, then the univariate forecasting equation with the GPDR is a reparameterization of the forecasting regression using the PDR. It follows then that the predicted values of the *ERP* are identical to those obtained by regressing the *ERP* on the log-transformed PDR:

$$pdr_t = p_t - d_t. \tag{19}$$

This approach of time-invariant $\hat{\alpha}$ and $\hat{\beta}$ assumes that there is a stable cointegration relationship between log-prices and log-dividends, which is implausible because of the time-varying economic conditions and market composition.

Secondly, if we restrict only $\hat{\alpha}_t$ to be time-invariant and use $\hat{\beta}_t$ as estimated for the INC-GPDR, then the predicted values of the *ERP* are identical to those obtained by regressing the *ERP* on the log-transformed MPDR, as given by:

$$\widehat{mpdr}_t = p_t - \hat{\beta}_t d_t. \tag{20}$$

Polimenis and Neokosmidis (2016) interpret the logarithmic MPDR as the mean-reverting error of a long-run equilibrium between log-prices and log-dividends, and interpret the recursive estimated β_t to be the estimation of a static population slope coefficient β. By using recursive estimations of β_t, Polimenis and Neokosmidis (2016) do not solve the two failures of the classical price-dividend ratio that we raise in this paper. The assumption of $\hat{\alpha}_t$ to be time-invariant affects the predictions, whenever a different estimation for $\hat{\alpha}_t$ is used for the INC-GPDR. Our approach is also fundamentally different by allowing β_t to be a possibly rapidly time-varying elasticity coefficient. Reflecting the evidence of Chen et al. (2016), we allow for the price-dividend ratio to revert to its average at different speeds, accounting for all the time-varying variables that may affect and disturb this relationship.

The inclusion of only recent (relevant) information should make the SMA-GPDR and the EWMA-GPDR more suitable for short- and medium-term *ERP* forecasting, whereas the staler MPDR is originally used for forecasting the *ERP* on a long horizon (more than 3 years) by Polimenis and Neokosmidis (2016). The comparison of the MPDR and GPDR ratios for forecasting the *ERP* at horizons longer than one year is an interesting topic for further research. It is beyond the scope of the present research, which focusses on exploiting the data-evidence for forecasting the intra-year *ERP*.

2.6. The GPDR at the intersection of fundamental and technical analysis

In their review on data-driven tools for forecasting the *ERP*, Neely et al. (2014) distinguish between the use of technical indicators and economic fundamentals. The former use past price and volume patterns to identify price trends believed to persist into the future, while the latter are motivated by a fundamental valuation analysis. We consider the approach of forecasting the *ERP* using the GPDR as an equity valuation model that integrates both the technical indicator and fundamental valuation approach. It has its origins in fundamental valuation, because prices are considered jointly with dividends in a log-log regression that follows from the valuation model in Campbell and Shiller (1988). However, the data-driven calibration of the GPDR parameters using weighted averages of log-prices, log-dividends and their centred cross-products can be interpreted as a technical approach. An interesting topic for further research is to compare the GPDR-based forecasting of the *ERP* with the stand-alone technical indicator and economic fundamentals approaches reviewed in Neely et al. (2014).

3. Data and summary statistics

The dataset used in this paper consists of monthly observations on nominal prices and dividends for the value-weighted S&P 500 index, as available from Amit Goyal's website[7]. Dividends are computed as the rolling sum of the dividends of the last 12 months, and are not reinvested over the last year period. The data goes from January 1947 until December 2014. The monthly *ERP* is calculated as in Equation (1). The first nine years of the dataset are used to estimate the time-varying parameters of the first MPDR and GPDRs which are then available from January 1956. It follows that the first sample used to estimate the relationship between the next period's *ERP* and the current log-transformed value of the predictor

[7] It is the same data as described in Section 1 of Goyal and Welch (2008) but updated until the end of 2014. It is available at `http://www.hec.unil.ch/agoyal/docs/PredictorData2015.xlsx`. In particular, the risk-free rate is defined as the total return when investing in a 3-month US Treasury Bill during the horizon for which the equity premium is calculated. Aggregate stock returns are taken from the Center for Research in Security Press (CRSP), while aggregate dividends from 1947 until 1987 are taken from Robert Shiller's website (`http://www.econ.yale.edu/?shiller/data.htm`), and from 1988 until 2014, they are from the S&P Corporation.

consists of 20 years from January 1956 until December 1975 (240 observations). The corresponding out-of-sample evaluation sample consists of 39 years from January 1976 until December 2014 (468 observations). By dividing the data up like this we avoid using data prior to World War II[8] and do not include the Oil Shock in our out-of-sample analysis avoiding to give the predictors an unfair advantage[9]. It also respects Hansen and Timmermann's (2012) findings that the forecast evaluation period should be a relatively large proportion of the available data.

Table 1 shows the summary statistics of the logarithmic PDR. Although the correlation between the log-transformed PDR and the next month's *ERP* has the expected negative sign, it is not significant. Table 1 further shows the summary statistics of the logarithmic INC-GPDR, SMA-GPDRs and EWMA-GPDRs. Consider first the column with the mean and median values. Since the GPDR is computed as the residual of a regression with time varying parameters, its value is close to 0. Another consequence is that its variability is smaller than the PDR. We see that the generalization reduces the standard deviations by two thirds compared to the PDR and MPDR. Also, the range of each GPDR is smaller. The GPDRs are persistent, as can be seen from the high values of the autocorrelation (between 0.92 and 0.96), but less so than than the PDR and MPDR for which the autocorrelation is above 0.99.

The last column, showing the correlation between the monthly *ERP* and a lagged predictor, is the most important one, as it indicates the usefulness to predict the monthly *ERP*. We expect this correlation coefficient to be negative. In fact, recall that the GPDR is designed to capture the misvaluation of the market portfolio and, for large (positive or negative) values of the GPDR a correction is expected to follow in terms of a future *ERP* that is lower (resp. higher) than average when the GPDR is positive (resp. negative). This is confirmed in Table 1, where we see that the lagged logarithmic Kitchin SMA-GPDR and EWMA-GPDR have a significant negative correlation with the monthly *ERP*. For the lagged logarithmic Julgaler SMA-GPDR and EWMA-GPDR the negative correlation is not significant. The correlation seems to lower as the window length tends to shrink.

Figure 2 compares the logarithmic Kitchin SMA-GPDR with the logarithmic PDR. The

[8] Lewellen (2004) recommends using only data after World War II because the properties of stock prices are much different prior to 1945.

[9] Goyal and Welch (2008) show that most valuation ratios benefit from including the Oil Shock of 1973–1975 in the out-of-sample analysis and state that the most recent decades after the Oil Shock can help shed light on whether a model is likely to still perform well nowadays.

Table 1: Summary statistics of the end-of-month logarithmic classical, modified and generalized price-dividend ratios over the period January 1956 until December 2014.

Variable	Window	Mean	Median	Min	Max	Sdev.	$\hat{\rho}_1$	$\hat{\rho}_{X_t,\ ERP_{t+1}}$
pdr		3.56	3.50	2.75	4.52	0.39	0.993	-0.04
mpdr	E	3.00	2.98	2.20	3.81	0.33	0.991	0.00
INC-*gpdr*	E	0.04	0.06	-0.69	0.83	0.31	0.989	-0.02
SMA-*gpdr*	K	0.01	0.02	-0.56	0.25	0.12	0.935	-0.09**
	J	0.03	0.05	-0.57	0.43	0.15	0.956	-0.05
EWMA-*gpdr*	K	0.01	0.02	-0.52	0.22	0.10	0.921	-0.07^{*}
	J	0.02	0.04	-0.55	0.34	0.13	0.942	-0.05

Note: The sample goes from January 1956 until December 2014 and monthly observations are used. All variables are in logs. The window length indicates the estimation window used for estimation of the coefficients of the variables, where E stands for expanding window from inception, while K and J stand for rolling window of length equal to the Kitchin and Julgaler cycle. $\hat{\rho}_1$ denotes the autocorrelation of order 1. $\hat{\rho}_{X_t,\ ERP_{t+1}}$ denotes the correlation between the monthly equity premium (ERP_{t+1}) and the lagged predictor (X_t). The significance levels are defined as follows: $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

horizontal line is the average for both ratios. The GPDR crosses its mean more frequently than its classical counterpart. The PDR has a large upward spike which started in the mid 90s. It has only returned briefly during the start of the financial crisis in 2008. Among others, Fama and French (2002) use this observation to develop a model where the PDR can be mean-reverting but the mean is allowed to vary between different regimes. Because of the rolling window calibration, the SMA-GPDRs and EWMA-GPDRs should naturally take care of this. Moreover, these regime switching methods often work well ex-post, but have little value for forecasting returns as these regime switches are difficult to detect ex-ante (Lettau and Van Nieuwerburgh, 2008). The GPDR approach does not suffer from this problem.

Finally, we must compare the GPDRs with the MPDR. The best forecast performance of the logarithmic MPDR is obtained when estimating the cointegration coefficient (β_t) using an expanding window (using w^{INC} as weights), in line with the results published in Polimenis and Neokosmidis (2016). For this reason, we report throughout the paper the results of the MPDR obtained as such[10]. Note that the sole difference between the INC-GPDR and the MPDR is that the former includes the time-varying α_t and the latter does not. Their correlation with the monthly *ERP* is insignificant.

[10] Polimenis and Neokosmidis (2016) also defend the use of a full-sample coefficient following the work of Lettau and Ludvigson (2001). However, following the terms in Lettau and Van Nieuwerburgh (2008) we only want to use a pure out-of-sample test instead of a pseudo one and thus do not include the MPDR with a full-sample coefficient in this paper.

Figure 2: Time series of the monthly price-dividend ratio and generalized price-dividend ratio (based on the simple moving average with the window length of the Kitchin cycle), in natural logarithm, for the S&P 500 over the period January 1956 - December 2014.

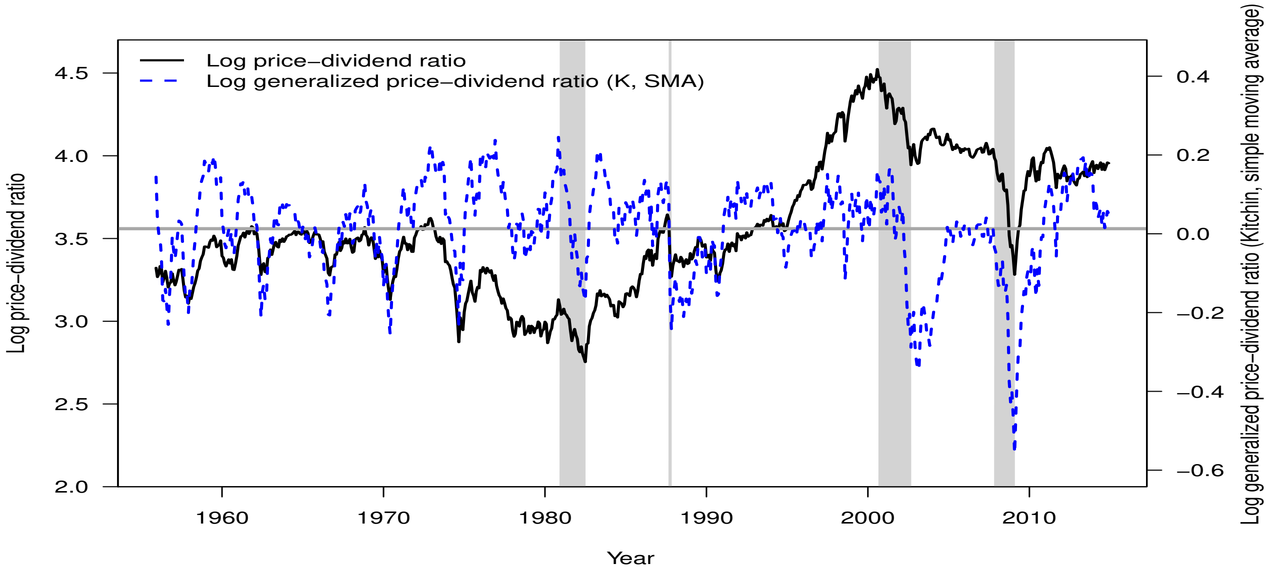

Note: The horizontal line is the average for both ratios. Shaded regions indicate the four largest drawdown periods of the S&P 500 during the out-of-sample period (01/1976 – 12/2014).

4. Out-of-sample evaluation of the forecasting precision

Forecasts of the *ERP* are useful for asset allocation and the valuation of risky assets, among others. In this section, we evaluate the forecasting accuracy of the monthly, quarterly, semi-annual and annual GPDR-based forecasts of the *ERP* using an out-of-sample statistical evaluation of forecast accuracy. We present full-sample statistics, and use graphical tools to display the time series evolution of the relative forecasting performance.

4.1. Benchmark models and evaluation criterion

Benchmark models. Goyal and Welch (2008) find that the prevailing historical average (HA) in (15) is a strong benchmark for out-of-sample comparison. It is often called the naive forecast as the historical average can be used to see if the predictor really adds some useful predictability in comparison with a naive investor who believes that the equity premium will be as it always has. Individual predictors typically fail to outperform the historical average in predicting the *ERP*. That is why the prevailing historical average is included as the main benchmark prediction. As alternative benchmark prediction methods, we also consider the random walk (RW) and an autoregressive of order 1 (AR(1)) prediction. These predictions are defined as follows:

$$\widehat{ERP}^{RW}_{t+1|t} = ERP_t, \tag{21}$$

$$\widehat{ERP}^{AR1}_{t+1|t} = \hat{c}_t + \hat{\rho}_t ERP_t, \tag{22}$$

with $\hat{c}_t$ and $\hat{\rho}_t$ being the OLS estimates of the AR(1) model parameters obtained from an expanding estimation window ending at time t.

Evaluation criterion. To measure the relative gains in forecasting precision, as compared to the naive forecast, we report the out-of-sample R–squared coefficient (R^2_{OOS}), which is defined as follows:

$$\mathrm{R}^2_{OOS} = 100 \times \left[1 - \frac{\sum_{s=1}^{n}(\widehat{ERP}_s - ERP_s)^2}{\sum_{s=1}^{n}(\widehat{ERP}^{\mathrm{HA}}_s - ERP_s)^2}\right], \tag{23}$$

where n is the number of out-of-sample forecasts, $\widehat{ERP}^{\mathrm{HA}}_s$ is the naive forecast obtained using Equation (15), and $\widehat{ERP}_s$ is the alternative ERP prediction obtained using the single-regime univariate, HAR, random walk or AR(1) model in Equation (11), (13), (21) and (22). Each alternative forecasting model is also combined with the switching forecasting equation in Equation (14). The R^2_{OOS} is multiplied by a 100 so it can be interpreted as a percentage. A positive value indicates a relative outperformance and vice versa in case of a negative value. According to Campbell and Thompson (2008), a monthly R^2_{OOS} near 0.5% already indicates an economically significant amount of predictability of the equity premium for an investor. The statistical significance level of the R^2_{OOS} is based on the bootstrap described in Goyal and Welch (2008). This bootstrap technique remains valid under the Stambaugh (1999) specification by preserving the autocorrelation structure of the predictor variable, and also preserves the cross-correlation structure of the two residuals by drawing in tandem (Rapach and Wohar, 2005).

4.2. Monthly equity premium

Table 2 shows the summary statistics of the predicted values for the monthly ERP obtained using the single-regime and switching forecasting equations for the PDR, MPDR, GPDRs, HAR-GPDR, RW, AR(1) and prevailing historical average. Note that the summary statistics of the out-of-sample values of the ERP are the same as for the RW prediction combined with the single-regime equation. By comparing the values for the predicted ERP with the values for RW, we can thus make the following conclusions. First of all, we see that all prediction methods have a lower mean value than the observed out-of-sample ERP and are thus too conservative. Secondly, because the *ERP* is predicted as a linear combination of the GPDR with slowly time-varying coefficients, it inherits the stability of the GPDR. The GPDR-based predictions using the single-regime forecasting equation are indeed more stable over time, as their standard deviation is often less than 10% of the standard deviation of the realized values of the *ERP*. This can also be seen in terms of the autocorrelation of order one, indicating the stability of the predicted values of the *ERP*. When the switching forecasting equation is used the autocorrelation and the standard deviation are lower and the average GPDR-based predictions move closer to the historical average. This is a result of the shifts to the prevailing historical average as a prediction, which on average happens half the time for all the predictors.

The two columns of Table 2 showing the R^2_{OOS} are the most important ones as they indicate the forecasting accuracy, relative to the prevailing historical average for forecasting the monthly *ERP*. The prevailing historical average seems to be the strongest benchmark as the alternative benchmark predictions are unable to outperform it when the single-regime forecasting equation is used. It is only when the switching forecasting equation is used, that the AR(1) is able to outperform the naive forecast but the accuracy gains are not statistically significant. If the single-regime forecasting equation is used only three predictors are able to significantly outperform the historical average, namely both Kitchin GPDRs and the HAR-GPDR. The Kitchin SMA-GPDR has the best results with a R^2_{OOS} of 0.53%. If the switching forecast equation is used the performance becomes even better, as it reaches a R^2_{OOS} of 1.52%. On top of that, the Julgaler GPDRs also significantly outperform the historical average. In general, the switching forecasting equation improves the performance of every predictor (alternative benchmarks included). The PDR, MPDR and INC-GPDR are never able to outperform the historical average.

The inclusion of α_t in the generalization seems to be important as the INC-GPDR is able to outperform the MPDR. Furthermore, the shorter estimation windows, which allow for more time-variation in the parameters, add extra forecasting power as every other GPDR is able to outperform the INC-GPDR. Estimating the time-varying parameters on the Kitchin cycle seems to work better than estimating them on the Julgaler cycle. A preliminary conclusion that can be drawn from this is that the forecasting precision is improved by discounting older information faster. This may also be the reason why for the Julgaler cycle, the EWMA-GPDR has better results than the SMA-GPDR. Curiously, the EWMA-GPDR performs worse than the SMA-GPDR when the Kitchin cycle is used. This may be caused by putting too much weight on recent observations which results in unstable estimates of the time-varying parameters. The R^2_{OOS} of the single-regime and switching HAR-GPDR, respectively 0.40% and 1.18%, are lower than the R^2_{OOS} of the single-regime and switching Kitchin SMA-GPDR. We can thus conclude that combining the information of the Kitchin and Julgaler cycle does not improve the forecasting accuracy of the GPDR. Based on these results, we see that the estimation of the time-varying parameters between logarithmic prices and dividends benefits the most of using the Kitchin cycle, especially when the simple moving average method is used.

In the remainder of the paper, we take the Kitchin SMA-GPDR as the default implementation of the GPDR and call it the GPDR when there is no confusion possible.

Table 2: Summary statistics of the predicted equity premium for the various forecasting methods considered, and its precision, as measured by the out-of-sample R^2 (relative to the prevailing historical average) are given for the forecasting of the one-month S&P 500 equity premium over the out-of-sample period January 1976 until December 2014, using the single-regime and the switching forecasting equations. We further report the frequency of naive forecasts used by the switching forecasting equation.

		Single-regime				*Switching*				
Predictor	Window	Mean (%)	Sdev. (%)	$\hat{\rho}_1$	R^2_{OOS}	*SW* (%)	Mean (%)	Sdev. (%)	$\hat{\rho}_1$	R^2_{OOS}
Naive	E	0.33	0.08	0.987	0.00					
RW		0.51	4.36	0.047	−90.29	50.43	0.38	3.31	0.134	−53.18
AR(1)	E	0.35	0.25	0.118	−0.19	50.21	0.33	0.20	0.319	0.11
pdr		0.26	0.34	0.985	−1.04	51.07	0.29	0.24	0.494	−1.16
mpdr	E	0.37	0.18	0.971	−0.53	46.37	0.35	0.14	0.809	−0.48
INC-*gpdr*	E	0.33	0.14	0.961	−0.38	50.00	0.33	0.20	0.319	−0.25
SMA-*gpdr*	K	0.40	0.48	0.947	0.53^{**}	52.99	0.37	0.32	0.653	1.52^{***}
	J	0.33	0.31	0.980	−0.06	51.71	0.34	0.22	0.719	0.38^{**}
EWMA-*gpdr*	K	0.38	0.33	0.934	0.28^{*}	52.78	0.36	0.21	0.641	0.72^{***}
	J	0.34	0.29	0.969	−0.02	51.07	0.34	0.20	0.702	0.38^{**}
HAR-*gpdr*	K+J	0.39	0.47	0.950	0.40^{**}	51.28	0.36	0.31	0.645	1.18^{***}

Note: This table contains the summary statistics of the predictions of the monthly *ERP*, the out-of-sample R^2 (R^2_{OOS}) for the single-regime and switching forecasting equations and the frequency of naive forecasts used by the switching forecasting equation. The R^2_{OOS} is calculated with respect to the prevailing historical average and is multiplied by a 100 so it can be interpreted as a percentage (see Equation (23)). The significance levels for the R^2_{OOS} are based on a bootstrap with a (one-sided) Null hypothesis of the naive forecast not performing inferior than the alternative model. The mean and standard deviation of the univariate predictions are multiplied by a 100 so they can be interpreted as a percentage. *SW* is the frequency of naive forecasts used by the switching forecasting equation in a percentage of the total amount of forecasts. The window length indicates the estimation window used for estimation of the coefficients of the variables, where E stands for expanding window from inception, while K and J stand for rolling window of length equal to the Kitchin and Julgaler cycle. The significance levels are defined as follows: $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$. The out-of-sample period contains 468 monthly observations. The initial regression is based on 240 monthly observations and an expanding estimation window is used.

4.3. Graphical analysis

Goyal and Welch (2003, 2008) propose to use a graphical approach analyzing the net-difference of the cumulative sum-squared error (CSSE) between the naive forecast and the more sophisticated prediction, which is in our case the GPDR-based prediction of the *ERP*:

$$CSSE_t = \sum_{s=1}^{t} (\widehat{ERP}_s^{\text{HA}} - ERP_s)^2 - (\widehat{ERP}_s^{\text{GPDR}} - ERP_s)^2, \tag{24}$$

and similarly for other prediction methods. Ideally, the out-of-sample CSSE performance of the GPDR must satisfy three criteria. Firstly, the model's performance needs a general (if not irregular) upward drift. Secondly, the drift should not be in a short or unusual sample period (this is why we left the Oil Shock out of the out-of-sample analysis). At last, the drift should remain positive over the most recent several decades, otherwise it may seem that the model has lost its forecasting power.

We evaluate these conditions in Figure 3, where we plot the time series of CSSE in (24) for the PDR, MPDR and the Kitchin SMA-GPDR. Similar plots are obtained when the Kitchin EWMA-GPDR and HAR-GPDR are used, as already indicated in Table 2. The upper (a) panel concerns the single-regime univariate forecasting equation, whereas the lower (b) panel shows the graphical result of the switching forecasting equation. The general interpretation of the figure is that the line represents the difference in CSSE between the naive model and the alternative model. An upward slope means that the alternative model is performing better than the prevailing historical average and vice versa. If the graph is in the region above zero it means a cumulative outperformance (and positive R^2_{OOS}) with respect to the prevailing historical average. It is easy to see how changing the starting or ending date would impact the conclusion of the out-of-sample result by shifting the horizontal line. The shaded areas indicate the four largest drawdowns of the S&P 500 during the out-of-sample period.

Figure 3: Out-of-sample cumulative sum-squared error (CSSE) performance of the prevailing historical average-based *ERP* against the use of alternative *ERP* model predictions.

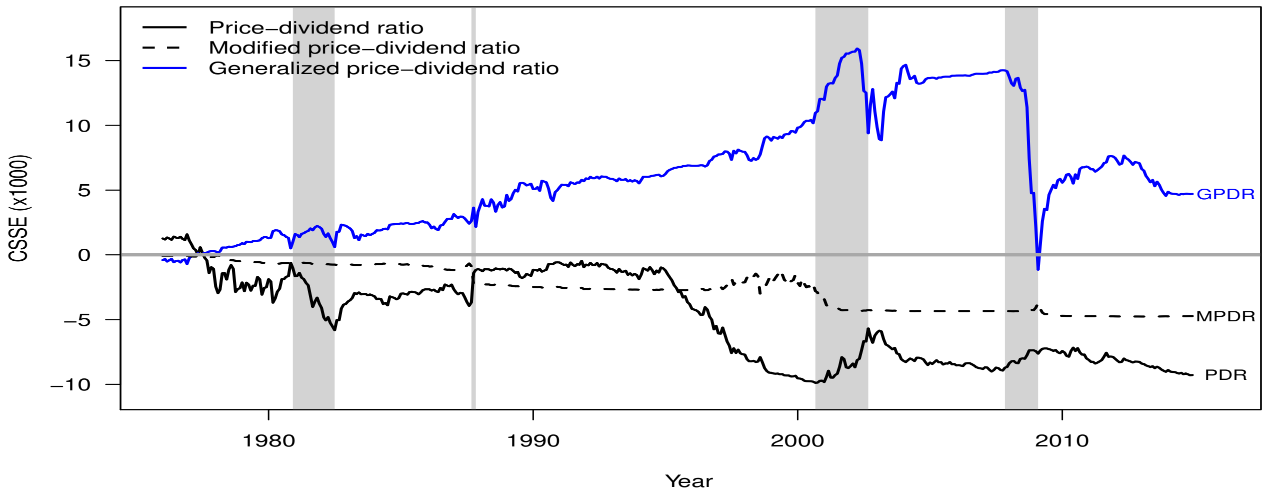

(a) Single-regime univariate forecasting model for the *ERP*

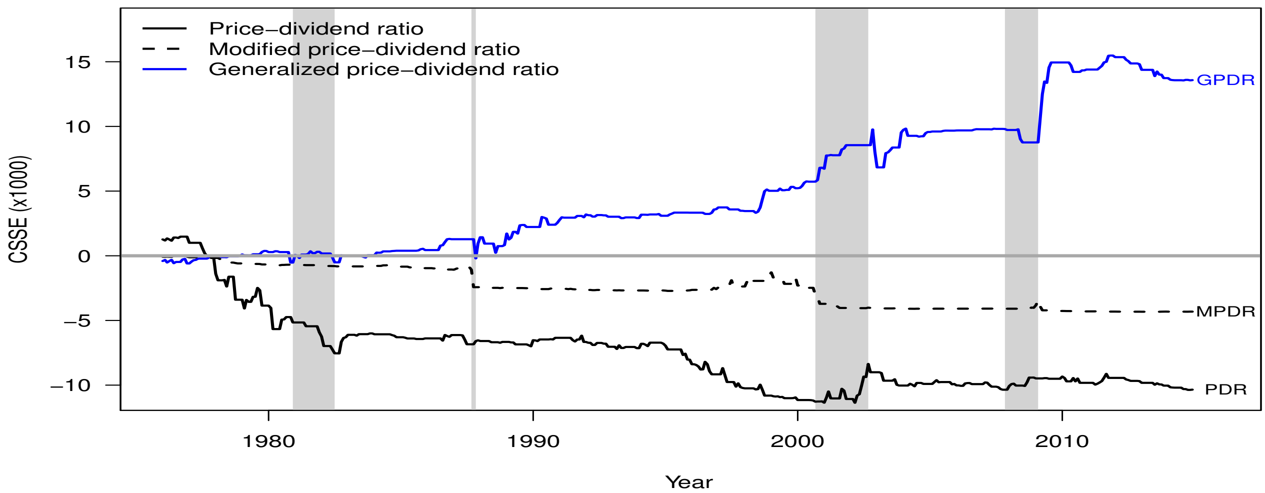

(b) Switching forecasting model for the *ERP*

Note: Time series plot of the cumulative sum-squared error of the prevailing historical average minus the cumulative sum-squared error of the predictor mentioned for forecasting the S&P 500 equity premium on a monthly forecasting horizon from January 1976 until December 2014 (x1000). Shaded regions indicate the four largest drawdown periods of the S&P 500 during the out-of-sample period.

The GPDR is almost always cumulatively outperforming its classical and modified counterpart as well as the naive forecast. Figure 3a shows a steady positive slope with some irregularities that mostly happen during the four indicated drawdown periods. In the first three drawdown periods the GPDR first performs better than the prevailing historical average, only to lose this outperformance in the second part of the drawdown period. The relative performance is the same at the beginning as at the ending of the drawdown period. However, the same does not hold for the last drawdown (i.e. the financial crisis 2007–2009). The GPDR in combination with the single-regime univariate forecasting equation loses all its outperformance during this period. This last major drawdown is the only reason why the switching forecasting equation is outperforming the single-regime univariate forecasting equation. Using the switching forecasting equation combined with the GPDR makes those declines disappear, but also causes the slope to become flat more often (52.99% of the predictions). This occurs when the data indicates that a more unpredictable market period has been reached and that an investor is better off using the naive forecast instead of the alternative model. Overall, Figure 3b confirms that the switching forecasting equation augments the statistical performance of the GPDR compared to the historical average.

4.4. Quarterly, semi-annual and annual equity premium

The previous analysis has shown that the GPDR leads to more accurate forecasts of the monthly equity premium. We now show that the GPDR has also added value in terms of accuracy gains when forecasting the quarterly, semi-annual and annual equity premium. Increasing the horizon makes the predictions more useful for tactical asset allocation with less frequent rebalancing moments, the valuation of equity investments or strategic asset allocation.

Table 3 shows the R^2_{OOS} for the PDR, MPDR, Kitchin SMA-GPDR and the alternative benchmark models relative to the historical prevailing average for forecasting the trailing 3, 6 and 12-month ERP, using the single-regime and switching forecasting equations. It further shows the frequency of naive forecasts used by the switching forecasting equation in a percentage of the total amount of forecasts (SW). To account for the longer return length and the overlapping nature of the returns, we add the return length to the window lengths (S) of the state variable in Equation (16) (i.e. for the trailing 3-month ERP the state variable uses window lengths 6 and 9 instead of 3 and 6).

The R^2_{OOS} results in Table 3 show that the prevailing historical average remains the strongest

Table 3: The out-of-sample R^2 (relative to the prevailing historical average) is given for the forecasting of the trailing 3, 6 and 12 month S&P 500 equity premium over the out-of-sample period January 1976 until December 2014 for the PDR, MPDR, Kitchin SMA-GPDR and the alternative benchmark models, using the single-regime and switching forecasting equation, together with the frequency of naive forecasts used by the switching forecasting equation.

	Single-regime			*Switching*					
	R^2_{OOS}			R^2_{OOS}			*SW* (%)		
Predictor	3	6	12	3	6	12	3	6	12
RW	-89.82	-91.06	-99.06	-31.55	-61.67	-56.85	44.66	36.11	36.54
AR(1)	-41.89	-65.23	-86.85	-14.78	-45.47	-49.23	46.15	36.97	36.75
pdr	-3.39	-8.97	-22.12	-1.81	-1.32	-7.96	52.78	51.92	48.72
mpdr	-1.44	-2.29	-5.45	-1.21	-0.55	-0.90	43.80	48.72	44.66
SMA-*gpdr*	1.63***	3.45***	3.34***	3.04***	3.38***	5.77***	54.27	58.55	56.20

Note: This table contains the out-of-sample R^2 (R^2_{OOS}) for the single-regime and switching forecasting equations and the frequency of naive forecasts used by the switching forecasting equation. The R^2_{OOS} is calculated with respect to the prevailing historical average and is multiplied by a 100 so it can be interpreted as a percentage (see Equation (23)). The significance levels for the R^2_{OOS} are based on a bootstrap with a (one-sided) Null hypothesis of the naive forecast not performing inferior than the alternative model. *SW* is the frequency of naive forecasts used by the switching forecasting equation in a percentage of the total amount of forecasts. The significance levels are defined as follows: $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$. The out-of-sample length contains 468 monthly observations. The initial regression is based on 240 monthly observations and an expanding estimation window is used.

benchmark for all the considered longer horizons as the RW and AR(1) are unable to outperform it, even combined with the switching forecasting equation. We further see that the single-regime and switching PDR and MPDR forecasting equations are also unable to outperform the historical average at any horizon. The R^2_{OOS} of the GPDR varies between 1.63% and 5.77% and is statistically significant at all horizons for both the single-regime and switching forecasting equations, indicating that the relative forecasting outperformance holds at longer horizons. In general, the switching forecasting equation improves the forecasting accuracy of each predictor.

When we analyze the frequency of naive forecasts used by the switching forecasting equation in Table 3, we see that the RW and AR(1) use the naive forecast less often for longer forecasting horizons, declining from about 50% to 36%, while the GPDR uses it slightly more often.

The bottom line is that the GPDR succeeds in significantly outperforming the prevailing historical average, while all the considered benchmarks are unable to, and that the switching forecasting equation improves the forecasting accuracy.

5. Out-of-sample evaluation of the economic value for tactical asset allocation

Our empirical results suggest that the GPDR leads to more accurate out-of-sample forecasts of the *ERP* at the intra-year horizon. We now examine whether these statistical gains can also have economic value when a mean-variance utility maximizing investor uses the GPDR-based prediction of the *ERP* for tactical asset allocation.

5.1. Tactical asset allocation implementation

We consider a risk-averse investor with mean-variance preferences who allocates his wealth across equities and risk-free bills at regular rebalancing times, namely at the last trading day of each month, quarter and year. The optimal fraction of his wealth that he allocates to equities (ω_t) is determined via the following objective function:

$$\max_{\omega_t} \mathbb{E}_t[r^p_{t+1}] - \frac{1}{2}\gamma \text{Var}_t(r^p_{t+1}), \tag{25}$$

where $r^p_{t+1} = r^f_{t+1} + \omega_t(r^m_{t+1} - r^f_{t+1})$ and, conditional on the information available at time t, $\mathbb{E}_t[r^p_{t+1}]$ is the expected value of the portfolio return and $\text{Var}_t(r^p_{t+1})$ is the corresponding variance of the portfolio returns[11]. The degree of relative risk aversion is denoted by γ. Following Goyal and Welch (2008), we set $\gamma = 3$, which is low enough to favour equity-tilted strategies and thus find an effect of the method used to predict the equity premium.

To solve the maximization problem in equation (25) we consider a myopic investor with a one-period investing horizon. From the first order conditions, it follows directly that the optimal allocation weight to equities during the month $t+1$ at the end of every month t is given by:

$$\omega_t = \frac{\mathbb{E}_t[r^m_{t+1} - r^f_{t+1}]}{\gamma \text{Var}_t(r^m_{t+1})}, \tag{26}$$

where $\mathbb{E}_t[r^m_{t+1} - r^f_{t+1}]$ is a prediction of the *ERP* over the investment horizon, and $\text{Var}_t(r^m_{t+1})$ is a prediction of its the variance. Both predictions are for time period $t+1$ and are conditional on the information available in time period t. We use the prevailing historical average, and the single-regime univariate and switching forecasting equations of the PDR, the MPDR and the Kitchin SMA-GPDR for the prediction of the *ERP*. We further assume the investor to use a 5-year moving window of monthly stock returns to estimate the variance. It is important to note that the obtained results are only for comparison and cannot be considered to be the

[11] We refer to Ardia et al. (2016) and references therein for a more general tactical portfolio optimization problem that involves also optimizing the allocation of the equity portfolio to the underlying stocks.

optimal results with these *ERP* predictions as ω_t also depends on the forecast of the variance. In order to safeguard the analysis against extreme estimation errors, we use in practice a constrained version of the optimized portfolio, obtained by imposing the box constraints that $0 \leq \omega_t \leq 1.5$ and thus allowing for a leverage of at most 50% (Campbell and Thompson, 2008).

In Table 4 we report for each portfolio its average equity weight (ω_t), and its annualized portfolio return, volatility, downside deviation, Sharpe ratio, Sortino ratio, turnover and certainty equivalent return gain. The mean portfolio return is geometrically annualized. The downside deviation aims to make a distinction between "good" and "bad" volatility and to account for the asymmetric distribution of returns. It is estimated as the square root of the mean squared deviation of the returns and risk free rate, when the returns are below the risk-free rate. It thus only includes the negative, or bad, volatility. The Sharpe ratio is the geometrically annualized excess mean portfolio return divided by its corresponding standard deviation. The Sortino ratio is similar to the Sharpe ratio but uses the downside deviation instead of the standard deviation. To show the potential impact of transaction costs, we further report the annualized turnover. It is the average wealth traded (i.e. bought and sold) per year in percentage of the total portfolio size (DeMiguel et al., 2009).

Finally, we also estimate the certainty equivalent return (CER) for each portfolio, defined as the return a risk-free investment should offer in order to yield the same expected utility as the risky portfolios that we consider. The estimation is based on equalizing the out-of-sample mean-variance utility, and thus leads to a CER estimate equal to:

$$CER_p = \hat{\mu}_p - \frac{1}{2}\gamma\hat{\sigma}_p^2, \tag{27}$$

where $\hat{\mu}_p$ is the arithmetic average of the investor's realized portfolio returns over the whole forecasting period and $\hat{\sigma}_p^2$ is the portfolio variance. To facilitate the interpretation, we follow Neely et al. (2014) by reporting the CER gain of the portfolios using the *ERP* predictions based on the PDR, MPDR and GPDR compared to the portfolio using the prevailing historical average. We multiply the CER gain by 100 to interpret it as a percentage. We further multiply it by 12 for the monthly allocation exercise and by 4 when the rebalancing is at a quarterly frequency in order to annualize it. This allows us to interpret the resulting annualized CER gain as the annual percentage fee that a naive investor would be willing to pay in order to use the alternative prediction method instead of the prevailing historical average.

5.2. *Out-of-sample results for tactical asset allocation*

We compare portfolios based on monthly, quarterly and annual *ERP* forecasts obtained using the prevailing historical average, and the single-regime univariate and switching forecasting equations of the PDR, the MPDR and the Kitchin SMA-GPDR. The performance results of the dynamic tactical asset allocation exercise can be found in Table 4. The single-regime PDR-based portfolios have the lowest average equity weights at about 60%, while the single-regime MPDR and historical average based portfolios have the highest at about 80%. The average equity weight of the single-regime GPDR is in between those at about 75%. Using the switching forecast equation generally increases the average equity weight which is caused by the frequent use of the naive forecast.

We see that the GPDR-based portfolios always generate the highest annualized mean portfolio return, both when the single-regime and switching forecasts equations are used. This comes at the price of a slightly higher volatility, but the annualized Sharpe ratio indicates that the excess return over the other portfolios compensates the excess volatility. Inspecting the downside deviation and Sortino ratio, we find that the excess volatility that is created using the GPDR-based predictions is mainly so-called "good" volatility, meaning that most of the excess volatility is from portfolio returns higher than the risk-free rate.

The relative outperformance in portfolio measures is confirmed by the annualized CER gain which is always positive for the GPDR-based portfolios. A naive investor would be willing to pay an annual percentage fee varying between 0.85% and 3.28%. For the monthly and quarterly rebalancing the single-regime GPDR-based portfolios perform better than the switching portfolios, while it is the other way around for annual rebalancing. All the considered benchmarks are unable to outperform the prevailing historical average, except for the portfolio using the **annual** predictions based on the PDR switching forecast equation.

The results above are for gross returns and thus before transaction costs. In Table 4 we see that the annualized turnover of the best performing portfolios, namely the single-regime GPDR-based portfolios for monthly and quarterly rebalancing and the switching GPDR-based portfolio for annual rebalancing, is a lot higher compared to the prevailing historical average and PDR-based portfolios. The outperformance in terms of net returns thus depends on the level of transaction costs. We verified that for proportional transaction costs up to 100 basis points the mean-variance utility maximizing investor is better off using the GPDR predictions compared to all the considered benchmarks. Moreover, when imposing

Table 4: The average ω_t and annualized return, volatility, downside deviation, Sharpe ratio, Sortino ratio, turnover and certainty equivalent return gain (relative to the prevailing historical average) of a portfolio using the equity premium forecasting approach compared to the portfolio obtained using the naive equity premium forecast when the portfolio is rebalanced on a monthly, quarterly and annual basis for the forecasting period of January 1976 until December 2014.

	Single-regime				*Switching*		
	Naive	*pdr*	*mpdr*	SMA-*gpdr*	*pdr*	*mpdr*	SMA-*gpdr*
Panel A: Monthly rebalancing							
Average Equity Weight (%)	77.28	60.63	78.52	75.13	67.10	77.92	78.08
Ann. Return (%)	9.40	6.49	8.62	13.29	7.09	8.62	11.97
Ann. Volatility (%)	13.05	13.06	13.91	15.30	12.69	13.57	14.26
Ann. Downside Deviation (%)	10.01	9.55	10.79	10.48	9.71	10.68	10.13
Ann. Sharpe Ratio	0.32	0.11	0.25	0.52	0.15	0.25	0.46
Ann. Sortino Ratio	0.42	0.15	0.32	0.75	0.20	0.32	0.65
Ann. Turnover (%)	26.74	71.09	41.61	176.05	268.27	62.02	228.76
Ann. CER gain (%)	0.00	-2.74	-0.94	2.87	-2.07	-0.85	2.00
Panel B: Quarterly rebalancing							
Average Equity Weight (%)	78.37	60.88	78.95	75.48	70.03	78.55	78.35
Ann. Return (%)	9.42	6.80	8.33	13.70	8.11	8.73	11.20
Ann. Volatility (%)	14.58	13.97	15.20	16.61	14.35	14.98	15.83
Ann. Downside Deviation (%)	10.43	9.75	11.37	10.13	10.14	11.16	10.42
Ann. Sharpe Ratio	0.29	0.12	0.21	0.50	0.21	0.24	0.37
Ann. Sortino Ratio	0.40	0.17	0.28	0.81	0.29	0.32	0.56
Ann. Turnover (%)	19.70	43.93	28.23	112.02	142.97	34.86	114.89
Ann. CER gain (%)	0.00	-2.30	-1.19	3.28	-1.17	-0.75	1.25
Panel C: Annual rebalancing							
Average Equity Weight (%)	80.12	61.59	77.35	72.94	77.06	81.78	81.33
Ann. Return (%)	9.67	7.81	8.62	11.12	10.14	9.39	11.27
Ann. Volatility (%)	17.67	14.73	17.86	19.27	16.08	17.67	17.63
Ann. Downside Deviation (%)	11.78	8.70	12.56	12.01	8.72	12.04	10.93
Ann. Sharpe Ratio	0.25	0.19	0.19	0.31	0.30	0.24	0.34
Ann. Sortino Ratio	0.38	0.30	0.27	0.48	0.56	0.35	0.54
Ann. Turnover (%)	14.26	23.36	20.29	46.64	39.24	18.55	36.98
Ann. CER gain (%)	0.00	-1.10	-1.06	0.85	0.78	-0.27	1.57

Note: This table shows the average equity weight (ω_t), the out-of-sample annualized return (%), volatility (%), downside deviation (%), Sharpe ratio, Sortino ratio, turnover (%) and certainty equivalent return gain with respect to the prevailing historical average (%) for each portfolio. Portfolios are rebalanced on a monthly, quarterly and annual basis. The relative risk aversion coefficient (γ) is 3. The window length indicates the estimation window used for estimation of the coefficients of the variables. The out-of-sample length contains 468 monthly observations. The initial regression is based on 240 monthly observations and an expanding estimation window is used.

more strict box constraints such as $0 \leq \omega_t \leq 1$ (no leverage allowed) or $0.4 \leq \omega_t \leq 0.8$ (more conservative constraints), the turnover of the portfolio can be substantially reduced while the relative outperformance remains.

In conclusion, the use of the GPDR-based *ERP* forecasts substantially improves the out-of-sample performance of tactical allocation portfolios. This result is robust to reasonable values of transaction costs and holds for monthly, quarterly and annually rebalancing of the portfolio.

6. Generalized financial ratios

Whenever a financial ratio compares the current price of the equity portfolio with a characteristic of that portfolio, it may benefit from our proposed generalization. Our contribution thus goes beyond generalizing the price-dividend ratio (**PDR**), and its use in forecasting the equity premium. The same economic intuition that is founding the generalization of the PDR can be used as a motivation to improve upon other popular price-based financial ratios, such as the (cyclically adjusted) price-earnings ratio, the price-earnings to growth ratio, the price-to-book ratio and the bond-equity yield ratios. In this section, we thus show that the framework of generalizing financial ratios that we propose is flexible and leads to the family of Generalized Financial Ratios (GFRs). We generalize financial ratios constructed out of other proxies and use them to predict the monthly *ERP*. Key results are shown and discussed.

6.1. Other choices for the fundamental characteristic

The current analysis shows that the aggregate dividends are a useful variable to predict the *ERP*. In this section we consider alternative predictive variables of three kinds. Firstly, earnings (E) variables related to the popular price-earnings ratio (PER), the cyclically adjusted price-earnings ratio (CAPE) and price-earnings to growth ratio (PEG). The CAPE uses real values instead of nominal ones. These real earnings are also a 10-year simple moving average to adjust for business cycles. The PEG resembles to the PER, but includes the growth rate (g) in the fundamental characteristic. Secondly, the book value of the S&P 500 firms (BV), as used in the price-to-book ratio (PB). Thirdly, the firm valuation as obtained under the constant dividend (earnings) perpetuity valuation model, with the inclusion of the long-term risk-free rate (i.e., D/Rf and E/Rf, respectively). The latter assumes that the dividends (earnings) remain constant at the last observed value. We thus obtain the Bond-Equity

Table 5: Financial ratios and their fundamental characteristics.

Predictor	*pdr*	*per*	*cape*	*peg*	*pb*	*bedyr*	*beeyr*
X_t	D_t	E_t	CA E_t	E_t/g_t	BV_t	D_t/Rf_t	E_t/Rf_t

Note: X_t is the fundamental characteristic at time t. CA stands for cyclically adjusted and means the 10-year simple moving average of the real values. D_t and E_t are the S&P500 nominal dividends and earnings. BV_t is the S&P 500 book-value. g_t is the 10-year growth rate of nominal earnings. Rf_t is the 10-year US Treasury constant maturity rate.

Yield Ratio (BEYR), studied in Giot and Petitjean (2007), among others. Note that it can be seen as an extension of the PEG, but instead of including the growth rate in the fundamental characteristic the long-term risk-free rate (Rf) is included. A distinction is made between the Bond-Equity Dividend Yield Ratio (BEDYR) and the Bond-Equity Earnings Yield Ratio (BEEYR). All the alternative fundamental characteristics and their corresponding predictor can be found in Table 5.

To formalize the generalization, denote X_t as the fundamental characteristic used such that we can generalize equation (5) as follows:

$$GPXR_t(\lambda_t, \beta_t) = \frac{P_t}{\lambda_t X_t^{\beta_t}}. \tag{28}$$

The data used in this section is obtained from Amit Goyal's website[12] and from Robert Shiller's website[13]. Again, monthly observations are used. The period in this section is the same as in the previous ones. That goes for in- and out-of-sample. For the PER the nominal earnings, calculated as sum of the earnings of the last 12 months, are used. The CAPE uses real prices and earnings instead. The Consumer Price Index (CPI) is used to deflate the nominal values and to avoid a look-ahead bias 1947 is used as base year. The earnings growth rate is the 10-year growth rate of nominal earnings[14]. The book-value of the S&P 500 is based on the market-to-book ratio used in Goyal and Welch's (2008) paper. The long-term interest rate is the 10-year US Treasury constant maturity rate (GS10). For a more detailed description of the earnings and book-value we refer to Goyal and Welch (2008), and for the CPI and long-term interest rate to Robert Shiller's website.

[12] It is available at `http://www.hec.unil.ch/agoyal/docs/PredictorData2015.xlsx`.

[13] It is available at `http://www.econ.yale.edu/?shiller/data.htm`.

[14] The logarithmic PEG and generalized PEG require positive values for earnings growth. In practice, we impose this by winsorizing the estimated earnings growth numbers at 1%. This affects only 1.7% of the observations, which all occur during the last financial crisis (2007-2009).

6.2. *Out-of-sample results of the GFRs*

Table 6 shows the key results when other proxies than the dividends are used in the form of GFRs to predict the monthly *ERP*. The in-sample period goes from January 1956 until December 2014 and the out-of-sample period starts in January 1976. To generalize the ratios the Kitchin and the Julgaler cycle are used. Only the simple moving average estimation method is used as the previous results in this paper show that these are representative for the generalization. The single-regime univariate and switching forecasting equations are used. The results are all with respect to the prevailing historical average so they can be compared with the performance of all the previous predictors. We have an in-sample measurement in the form of an in-sample correlation, and the out-of-sample performance metrics R^2_{OOS} and the annualized CER gain.

The results of the generalization are less clear with the alternative fundamental characteristics. Sometimes it does help in statistical terms and not in economic terms and vice versa. The same can be concluded for using the switching forecasting equation. The generalized BEEYR is able to significantly outperform the prevailing historical average. The other variables are in general unable to outperform the prevailing historical average. The fact that it is hard to predict the monthly *ERP* with a single variable is in line with the findings of Goyal and Welch (2008). Compared to all these variables the SMA-GPDR based on the Kitchin cycle remains the best forecaster of the monthly *ERP*.

The bottom line of the investigation of using the generalized financial ratios is that the flexibility of the approach offers opportunities to use alternative stock characteristics than the dividends per share, but that, over the alternatives considered, the GPDR is the generalized financial ratio that we recommend to use, as it leads to the most reliable results in terms of statistical accuracy and economic benefits.

Table 6: The correlation between a lagged predictor and the monthly equity premium, the R^2_{OOS} (relative to the prevailing historical average) and the annualized certainty equivalence return gain (relative to the prevailing historical average) are shown for classic predictors and generalized predictors based on the simple moving average of the Kitchin and Julgaler cycle using the single-regime and switching forecasting equations.

	Single-regime						*Switching*					
Panel A: *Forecasting the monthly ERP using the classical financial ratios*												
Predictor	*per*	*cape*	*peg*	*pb*	*bedyr*	*beeyr*	*per*	*cape*	*peg*	*pb*	*bedyr*	*beeyr*
$\hat{\rho}_{X_t,\, ERP_{t+1}}$	-0.02	-0.03	-0.02	-0.01	-0.09^{**}	-0.07^{*}						
R^2_{OOS}	-0.60	-0.46	-1.00	-0.73	-0.59	-0.38	-0.20	-0.45	-0.66	-0.58	-1.43	-0.61
Ann. CER gain (%)	0.23	-1.51	-1.83	-1.31	-0.98	0.86	0.79	-1.47	-1.28	-1.24	-1.67	-0.04
Panel B: *Forecasting the monthly ERP using the generalized financial ratios*												
Predictor	*gper*	*cagpe*	*gpeg*	*gpb*	*gbedyr*	*gbeeyr*	*gper*	*cagpe*	*gpeg*	*gpb*	*gbedyr*	*gbeeyr*
Kitchin cycle (SMA)												
$\hat{\rho}_{X_t,\, ERP_{t+1}}$	-0.02	0.00	0.00	-0.01	-0.02	-0.08^{**}						
R^2_{OOS}	-0.40	-0.46	-0.71	-0.62	-0.52	0.58^{**}	-0.29	-0.36	-0.44	-0.40	-0.09	0.32^{**}
Ann. CER gain (%)	-0.94	-1.07	-0.80	-1.27	-0.72	1.02	-0.78	-0.91	-0.28	-0.72	-0.31	0.67
Julgaler cycle (SMA)												
$\hat{\rho}_{X_t,\, ERP_{t+1}}$	-0.02	0.00	-0.02	0.00	-0.06	-0.08^{**}						
R^2_{OOS}	-0.59	-0.56	-0.59	-0.98	0.07	0.63^{**}	-0.48	-0.32	-0.40	-0.62	0.06	0.32^{**}
Ann. CER gain (%)	-0.58	-1.04	-1.35	-1.29	0.90	1.59	-0.51	-0.76	-0.77	-0.67	0.41	0.58

Note: This table contains the correlation between a lagged predictor and the monthly equity premium, the R^2_{OOS} with respect to the prevailing historical average and the annualized certainty equivalence return gain with respect to the prevailing historical average (in %). $\hat{\rho}_{X_t,\, ERP_{t+1}}$ denotes the correlation between the monthly equity premium (ERP_{t+1}) and the lagged predictor (X_t). The in-sample correlation is based on 708 monthly observations from 01/1956 until 12/2014. The R^2_{OOS} is calculated with respect to the naive forecast and is multiplied by a 100 so it can be interpreted as a percentage (see Equation (23)). The significance levels for the R^2_{OOS} are based on the bootstrap with a (one-sided) Null hypothesis of the naive forecast not performing inferior than the alternative model. The out-of-sample length contains 468 monthly observations. The initial regression is based on 240 monthly observations and an expanding estimation window is used. The relative risk aversion coefficient (γ) is 3. The window length indicates the estimation window used for estimation of the coefficients of the variables. All the predictors are in logs. The significance levels are defined as follows: $^{***}p < 0.01$, $^{**}p < 0.05$, $^{*}p < 0.1$.

7. Conclusion

The price-dividend ratio (PDR) expresses the stock's price per unit of dividend the firm pays to its shareholders. There is only weak empirical support for the modern finance prediction that the PDR ought to predict either stock returns or dividend growth. We argue that predicting the equity premium using the information in the PDR can be substantially improved by accounting for the time-variation in the discount factor, the market composition, dividend policy and the investors' attitudes towards dividends and taxes. More specifically, we propose to compare the price per share with a time-varying function of the dividends per share. The time-variation in the parameters is intended to account for the aforementioned confounding factors causing the disconnection between the price-dividend ratio and the equity premium. We find that the resulting Generalized Price-Dividend Ratio (GPDR) improves the forecasting precision for the US intra-year equity premium from January 1976 until December 2014. Importantly, we document that these statistical gains also have economic value in terms of a better investment performance when using these predictions for tactical asset allocation. Similar gains in performance are found when generalizing financial ratios comparing the market valuation with corporate earnings or book value. Of all ratios considered, the **GPDR** performs best in predicting the intra-year US equity premium. The generalized bond-equity earnings yield ratio and the generalized bond-equity dividend yield ratio are second and third best. We obtain these results for the US aggregate stock market. In further research, one could follow Lawrenz and Zorn (2017) and explore whether the GPDR is able to predict (excess) stock returns in an international time-series or cross-sectional context.

References

Al-Malkawi, H.-A. N., Bhatti, M. I., Magableh, S. I., 2014. On the dividend smoothing, signaling and the global financial crisis. Economic Modelling 42, 159–165.

Ardia, D., Boudt, K., Wauters, M., 2016. The economic benefits of market timing the style allocation of characteristic-based portfolios. North American Journal of Economics and Finance 37, 38 – 62.

Bogle, J. C., 2005. Foreword. The intelligent investor: The classic text on value investing. By Benjamin Graham. HarperCollins Publishers Inc.

Brennan, M., Schwartz, E. S., Lagnado, R., 1997. Strategic asset allocation. Journal of Economic Dynamics and Control 21 (8–9), 1377–1403.

Campbell, J. Y., Shiller, R. J., 1988. The dividend-price ratio and expectations of future dividends and discount factors. Review of Financial Studies 1 (3), 195–228.

Campbell, J. Y., Thompson, S. B., 2008. Predicting excess stock returns out of sample: Can anything beat the historical average? Review of Financial Studies 21 (4), 1509–1531.

Chan, K., Karolyi, G., Stulz, R., 1992. Global financial markets and the risk premium on U.S. equity. Journal of Financial Economics 32 (2), 137–167.

Chen, S.-W., Hsu, C.-S., Xie, Z., 2016. Are there periodically collapsing bubbles in the stock markets? New international evidence. Economic Modelling 52, 442–451.

Cochrane, J. H., 1997. Where is the market going? Uncertain facts and novel theories. Economic Perspectives 21 (6), 3–37.

Cochrane, J. H., 2008. The dog that did not bark: A defense of return predictability. Review of Financial Studies 21 (4), 1533–1575.

Cochrane, J. H., 2017. Macro-finance. Review of Finance 21 (3), 945–985.

Corsi, F., 2009. A simple approximate long-memory model of realized volatility. Journal of Financial Econometrics 7 (2), 174–196.

Damodaran, A., 2009. Equity risk premiums (ERP): Determinants, estimation and implications–a post-crisis update. Financial Markets, Institutions & Instruments 18 (5), 289–370.

DeMiguel, V., Garlappi, L., Uppal, R., 2009. Optimal versus naive diversification: How inefficient is the 1/n portfolio strategy? Review of Financial Studies 22 (5), 1915–1953.

Duarte, F., Rosa, C., 2015. The equity risk premium: A review of models. Federal Reserve Bank of New York Economic Policy Review 21 (2), 39–57.

Faber, M., 2007. A quantitative approach to tactical asset allocation. Journal of Wealth Management 9 (4), 69–79.

Fama, E. F., French, K. R., 1988. Dividend yields and expected stock returns. Journal of Financial Economics 22 (1), 3–25.

Fama, E. F., French, K. R., 2001. Disappearing dividend: Changing firm characteristics or lower propensity to pay? Journal of Applied Corporate Finance 14 (1), 67–79.

Fama, E. F., French, K. R., 2002. The equity premium. Journal of Finance 57 (2), 637–659.

Genre, V., Kenny, G., Meyler, A., Timmermann, A., 2013. Combining expert forecasts: Can anything beat the simple average? International Journal of Forecasting 29 (1), 108–121.

Giot, P., Petitjean, M., 2007. The information content of the bond-equity yield ratio: Better than a random walk? International Journal of Forecasting 23 (2), 289–305.

Gordon, M. J., 1959. Dividends, earnings, and stock prices. Review of Economics and Statistics 41 (2), 99–105.

Goyal, A., Welch, I., 2003. Predicting the equity premium with dividend ratios. Management Science 49 (5), 639–654.

Goyal, A., Welch, I., 2008. A comprehensive look at the empirical performance of equity premium prediction. Review of Financial Studies 21 (4), 1455–1508.

Graham, J. R., Harvey, C. R., 2005. The long-run equity risk premium. Finance Research Letters 2 (4), 185–194.

Gupta, R., Mampho, P., M., 2013. Macroeconomic variables and South African stock return predictability. Economic Modelling 30, 612–622.

Hansen, P., Timmermann, A., 2012. Choice of sample split in out-of- sample forecast evaluation. Working paper, European University Institute, Florence, Italy.

Hull, B., Qiao, X., 2017. A practitioner's defense of return predictability. Journal of Portfolio Management 43 (3), 60–76.

Jadevicius, A., Hutson, S., 2014. A "family of cycles" – Major and auxiliary business cycles. Journal of Property Investment & Finance 32 (3), 306–323.

Kaustia, M., Knüpfer, S., 2008. Do investors overweight personal experience? Evidence from IPO subscriptions. Journal of Finance 63 (6), 2679–2702.

Korniotis, G. M., Kumar, A., 2011. Do older investors make better investment decisions? Review of Economics and Statistics 93 (1), 244–265.

Kuhnen, C. M., Knutson, B., 2011. The influence of affect on beliefs, preferences, and financial decisions. Journal of Financial and Quantitative Analysis 46 (3), 605–626.

Lawrenz, J., Zorn, J., 2017. Predicting international stock returns with conditional price-to-fundamental ratios. Journal of Empirical Finance. Forthcoming.

LeRoy, S. F., 1973. Risk aversion and the martingale property of stock prices. International Economic Review 14 (2), 436–446.

Lettau, M., Ludvigson, S., 2001. Consumption, aggregate wealth, and expected stock returns. Journal of Finance 56 (3), 815–849.

Lettau, M., Van Nieuwerburgh, S., 2008. Reconciling the return predictability evidence. Review of Financial Studies 21 (4), 1607–1652.

Lewellen, J., 2004. Predicting returns with financial ratios. Journal of Financial Economics 74 (2), 209–235.

Lo, A., 2005. Reconciling efficient market with behavioral finance: The adaptive markets hypothesis. Journal of Investment Consulting 7 (2), 21–44.

Lucas, R. E. J., 1978. Asset prices in an exchange economy. Econometrica 46 (6), 1429–1445.

Nadaraya, E., 1964. On estimating regression. Theory of Probability & Its Applications 9 (1), 141–142.

Neely, C. J., Rapach, D. E., Tu, J., Zhou, G., 2014. Forecasting the equity risk premium: The role of technical indicators. Management Science 60 (7), 1772–1791.

Polimenis, V., Neokosmidis, I. M., 2016. The modified dividend–price ratio. International Review of Financial Analysis 45, 31–38.

Prat, G., 2013. Equity risk premium and time horizon: What do the U.S. secular data say? Economic Modelling 34, 76–88.

Rapach, D. E., Wohar, M. E., 2005. Valuation ratios and long-horizon stock price predictability. Journal of Applied Econometrics 20 (3), 327–344.

Rozeff, M. S., 1984. Dividend yields are equity risk premiums. Journal of Portfolio Management 12 (2), 68–75.

Shiller, R. J., 1984. Stock prices and social dynamics. Brookings Papers on Economic Activity 2, 457–498.

Stambaugh, R. F., 1999. Predictive regressions. Journal of Financial Economics 54 (3), 375–421.

Stock, J. H., Watson, M. W., 2004. Combination forecasts of output growth in a seven-country data set. Journal of Forecasting 23 (6), 405–430.

Watson, G., 1964. Smooth regression analysis. Sankhya: The Indian Journal of Statistics, Series A (1961-2002) 26 (4), 359–372.

Welch, I., 2000. Views of financial economists on the equity premium and on professional controversies. Journal of Business 73 (4), 501–537.