International Review of Financial Analysis 45 (2016) 31–38

Contents lists available at ScienceDirect

International Review of Financial Analysis

The modified dividend–price ratio

Vassilis Polimenis [a,*,1], Ioannis M. Neokosmidis [b]

[a] Business Division, Department of Economics, Aristotle University, 54124 Thessaloniki, Greece
[b] Department of Economics, Aristotle University, 54124 Thessaloniki, Greece

ARTICLE INFO

Article history:
Received 24 July 2015
Received in revised form 23 January 2016
Accepted 25 January 2016
Available online xxxx

JEL classification:
G11
G12
G13
G14
G17

Keywords:
Dividend–price ratio
Non-stationary ratios
Cointegrated dividend–prices
Modified dividend–price ratio
Out of sample forecasts
Return predictability

ABSTRACT

We show that log-dividends (d) and log-prices (p) are cointegrated, but, instead of de facto assuming the stationarity of the classical log dividend–price ratio, we allow the data to reveal the cointegration vector between d and p. We define the modified dividend–price ratio (mdp), as the long run trend deviation between d and p. Using S&P 500 data for the period 1926 to 2012, we show that mdp provides substantially improved forecasting results over the classical dp ratio. Out of sample, while the dp ratio cannot outperform the "simplistic forecast" benchmark for any useful horizon, an investor who employs the mdp ratio will do significantly better in forecasting 3-, 5- and 7-year returns with an R^2_{OS} of 7%, 26% and 31% respectively. In some sense mdp can be considered as a de-noising of the classical ratio as it addresses the major weakness in dp, its presumed inability in revealing business cycle variation in expected returns. Unlike dp, mdp exhibits positive correlation with the risk free return component, and can discern if a low dividend state coincides with a low yield state.

1. Introduction

The ability to forecast returns can easily be regarded as the most significant question for asset allocation, and one of the most important issues in the entire financial economics. After an early period where return predictability was approached with simplistic or brute-force methods, during the late 80s and early 90s, the literature proposed more sophisticated and smart ways to measure the ability of valuation ratios, and other statistics in predicting aggregate stock returns. Motivated mainly by practitioner views starting with the classic Graham and Dodd (1934) that high valuation ratios should carry positive information about future returns, Fama and French (1988) find that economically substantial return *predictability at a long horizon* exists. Long-horizon forecasts are the mechanical result of short horizon same-direction forecastability combined with a highly *persistent* forecasting variable. The persistence of a predictor variable leads to increased predictive slope coefficients for longer horizons.

Miller and Modigliani (1961) argued that dividend policy is irrelevant, and that stock prices should be driven by the "real" variable which is the earnings power of corporate assets. Yet, from early on dividend yields attained special importance as a forecasting variable due to the straightforward participation of the dividend yield in return formation, and its highly persistent dynamics which could provide predictability in long forecasting horizons via the mechanism outlined above.

Cochrane (1992, 2011) argues that for long horizons, long-run return and/or dividend growth predictability have to coincide with the variability of the log dividend–price ratio (dp)[2] Actually Cochrane (2011) goes one step further in arguing that (surprisingly) dp has no information about future dividend growth, and that almost all variation in dividend yields is driven by variation in discount rates. Powerful as it may be, this finding is based on two main assumptions, a) the *stationarity* of dividend yields and b) the assumed ability to recursively extend the Campbell and Shiller (1988) approximation to *infinity*.[3] Furthermore, there are some major problems with the predicting performance

* Corresponding author.
E-mail addresses: polimenis@econ.auth.gr (V. Polimenis), ineokosm@econ.auth.gr (I.M. Neokosmidis).
[1] Tel.: +30 2310 99 1317.

[2] In this paper lowercase letters always denote logs: $d_t = \log D_t$, $p_t = \log P_t$, and $r_t = \log R_t$.
[3] Engsted, Pedersen, and Tanggaard (2012) study the error of the Campbell–Shiller approximation in the presence of a non-stationary dp.

http://dx.doi.org/10.1016/j.irfa.2016.01.019
1057-5219/

Table 1
Summary Statistics.
We present the summary statistics for annual returns, equity premia, risk free rates, classic dividend–price ratio (dp) and our modified dividend–price ratio (mdp). The table shows the correlation matrix between the series as well as the mean, standard deviation and the autocorrelation coefficient based on AR(1) fitted model. Data are annual from 1926 to 2012.

	r_t	re_t	rf_t	dp_t	mdp_t	Mean	Std	AR(1)
r_t	1					0.09	0.20	0.06
re_t	0.99	1				0.06	0.20	0.05
rf_t	0.03	−0.12	1			0.04	0.03	0.93
dp_t	−0.25	−0.24	−0.05	1		−3.35	0.45	0.87
mdp_t	−0.34	−0.39	0.35	0.69	1	−2.05	0.26	0.70

of the dividend yield. Firstly, its weak performance in predicting returns and risk premia outside the sample used to determine the slope coefficient. Secondly, an inability in revealing high to medium frequency variation (i.e. business cycles) in expected returns and equity risk-premia. Over shorter than 7–10 year horizons, dividend–price ratios mainly predict themselves (Goyal & Welch, 2003). The poor Out-of-Sample (OS) performance of dividend–price ratio is exhibited in Goyal and Welch (2003); Welch and Goyal (2008) and Campbell and Thompson (2008).

1.1. The non-stationarity of the dividend yield

Econometrically, most researchers argue that dp is a stationary process based on *infinite* sample or *asymptotic* arguments, and take dp stationarity as a given assumption. But neither the data sets that we actually use, nor the time horizons that we use to evaluate our models' performance are infinite. At the same time, the majority of empirical studies on return predictability, cannot reject statistically (if not economically) the hypothesis of the presence of a unit root in the dividend–price ratio (Goyal & Welch, 2003; Lettau & Van Nieuwerburgh, 2008; Lettau & Ludvigson, 2005 among others).

We can see from summary statistics presented in Table 1, that the dividend–price ratio dp has an autocorrelation $\varphi = .87$. Clearly, this is a local alternative that unit root tests have not enough power to detect. Furthermore, it is known, as early as Kendall (1954) that typical estimation methods will tend to highly underestimate true persistence in finite samples.[4] In the following sections, we present robust econometric evidence against the stationarity of the classical dp. Not only is stationarity rejected via a straightforward ADF testing for dp, but using the more powerful test of a restriction on the cointegration vector for d and p we reject the hypothesis that log-dividends and log-prices are linked with a long run relationship of the form (d-p).[5] Econometrically, dp is at best a near non-stationary process.

Economically, the unquestionable requirement that stock prices cannot be far from corporate fundamentals for too long has often been interpreted in a strict sense that the log dividend–price ratio is stationary either in the full sample or at least in specific subsamples. The classical thinking about the behavior of dividend yield ratios is that dividends should represent a more or less "fixed" fraction of earnings, and earnings should represent a more or less "fixed" fraction of prices. Thus, most contemporary literature de facto assumes that the classical dp is a *stationary process* and should not include any trends. Generally speaking though, this is an economic requirement, which depends on a particular sample, rather than a hard fact. Actually, corporate officers have large discretion over payout options, and such discretion might impart unexpected structure into the dynamics of the dividend yield.

The fact that, over any finite period of time, dividends (and dividend growth) can be arbitrary, and delinked from asset prices, means we should neither be dogmatic about the time series properties of the dividend yield nor about its inability to predict dividend growth. Yet, generally speaking, both academia and practice have avoided tackling head-on the possibility of non-stationary dynamics in valuation ratios such as the dividend–price ratio, despite the fact that the hypothesis of a unit root in long horizon samples cannot be statistically rejected. The economical source of such non-stationarity in dividend yields is not easily understood. It could be the result of changes in dividend policy such as *dividend smoothing*, use of *share repurchases* in lieu of cash payments, or it could be induced by other changes of investors' attitudes toward dividends and taxes.

In any case, such changes in dividend policy will emerge in the data as a *slope differential* between dividends and prices. When we move away from dividend yield stationarity, assuming a deterministic long run equilibrium relation between dividends and prices is the next logical step still satisfying a "fundamentals" based asset pricing philosophy. In this paper we modify the dividend–price ratio by relaxing the stationarity assumption for the classical dp_t, and assuming a deterministic long run relation between dividends and prices; i.e. assume a cointegration vector of the form $d_t = \alpha + \beta p_t$, and allow the data to reveal the "true" cointegration vector $[1, -\beta]$.

In the above long-run relation, we define the modified dividend–price ratio as the stationary cointegration error of this long-run equilibrium, mdp = d–βp. We may then think of β as the unique population parameter that "fine tunes" dp by revealing the stationary trend deviation between dividends and prices. This modified ratio (mdp) is more informative than its non-stationary counterpart, the classical dp ratio. Effectively, in our analysis, the classical dp can be thought as the modified ratio, mdp, plus a (possibly) small I(1) noise term.

$$dp_t = mdp_t + (\beta - 1)p_t. \tag{1}$$

By not de facto assuming an *unreliable rejection* of the non-stationarity null for dp, the modified ratio presents a more reliable alternative, which allows for a richer representation of the d.g.p. Also, at $\varphi = .70$, mdp still has enough persistence in order to provide forecastability in long horizons. Before diving into a set of econometric tests, that will undoubtedly establish the superiority of using our trend-corrected modified dividend–price ratio in forecasting long-run returns, it is worth to first approach the *economic* ramifications of a non-stationary *dp* from a qualitative point of view.

In our setup, β provides the drift ratio between d and p. Roughly speaking, a $\beta < 1$ implies that dividends have been growing more slowly than prices. Having motivated the possibility for such a slope differential, and thus a non-stationary dp, the important question with respect to understanding the true dynamics of dp is whether such non-stationarity is only due to a deterministic time trend or it includes a unit root. The problem is that, as is now well understood, this question is inherently unanswerable for any finite sample (see Blough, 1992) since for any unit root process, and sample size T, there exists a stationary process that is indistinguishable. Another way to understand this issue is that the question of the inclusion of a unit root in the process is equivalent to finding whether the population spectrum at zero is zero or attains any positive value. This is clearly unanswerable, since in any sample there is no information about cycles of a period larger than the sample size. A realistic target for the financial economist should rather be to describe the data in a parsimonious way with low order autoregressions, since they are easier to estimate than high order moving average processes.

We show that an investor who employs the modified ratio (mdp) will improve his Out-of-Sample forecasting of 3-, 5- and 7-year returns with an R^2_{OS} of 7%, 26% and 31% respectively. Furthermore, an investor who has seen enough of the small (due to super-consistency) required early sub-sample to reliably infer population values for the cointegration coefficient between d and p, will actually improve his forecasts of the 5- and 7-year returns by an astonishing R^2_{OS} of 49%, and even attain a 3-year R^2_{OS} of 34%.

[4] Actually, even the Kendall bias correction for autocorrelation $-(1+3\varphi)/T$ is low.
[5] That the $[1, -1]$ vector spans the cointegration space.